5. PROJECT EASI/ED CANDIDATE ARCHITECTURES

This section defines four candidate framework architectures evaluated for Project EASI/ED. These architectures were selected from a collection of architectural alternatives (described in Appendix F) after careful consideration of the criteria described in Section 4. The candidate architectures are based on the process and data distribution strategies described in Subsection 2.2.

Subsection 5.2 identifies the goals, assumptions, and constraints that affect the candidate architectures. Subsection 5.3 describes the strategies used to populate each of the candidate architectures. Subsection 5.4 defines in detail the 4 candidate framework architectures for Project EASI/ED.

Appendix B contains descriptions of key technologies used to populate the candidate architecture models. This appendix also provides a representative list of commercially available solutions for the described technologies.

5.1 Project EASI/ED Candidate Architecture Overview

The candidate architectures for Project EASI/ED are:

- Centralized Processing/Centralized Data
- Distributed Processing/Centralized Data
- Distributed Processing/Replicated Data for Consolidation
- Distributed Processing/Replicated Data for Publication

This subsection introduces the major characteristics of each candidate architecture without consideration of Project EASI/ED specific configurations. Subsection 5.4 describes how each candidate framework architecture was configured to meet Project EASI/ED requirements for this evaluation.

5.1.1 Centralized Processing/Centralized Data

As illustrated in Figure 5-1, this architecture includes centralized data and application resources, which facilitate the execution of all data management, application, and presentation processing from a single computing resource. This architecture is the most secure and uses the least complicated data distribution strategy considered during this analysis. However, this strategy can adversely affect system scalability. It also introduces risk, as all system processes are dependent upon the availability of a single resource – the centralized database.

Figure 5-1 shows the Project EASI/ED residing on a single enterprise server and updating a single central database. The system communicates with students and trading partner systems over a network of public and private communication resources.

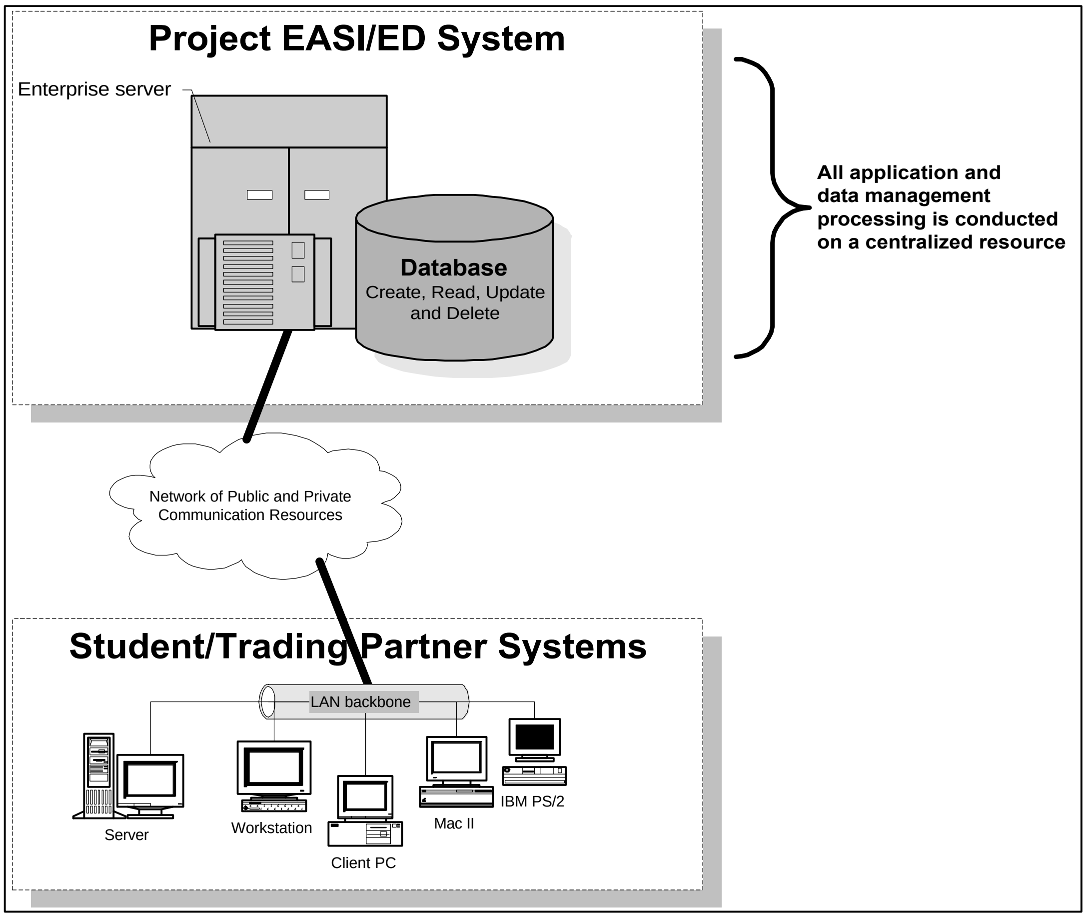

Figure 5-1. Centralized Processing/Centralized Data Architecture

5.1.2 Distributed Processing/Centralized Data

As illustrated in Figure 5-2, this architecture includes centralized data management resources, which facilitate the execution of all data management processing (transaction processing and decision support) from a single computing resource. However, unlike the fully centralized architecture, this architecture allows for the distribution of application and presentation processing resources and activities.

Figure 5-2 shows the Project EASI/ED processing distributed across multiple application processing servers. All data management processing is against a single central database. The system communicates with students and trading partner systems over a network of public and private communication resources.

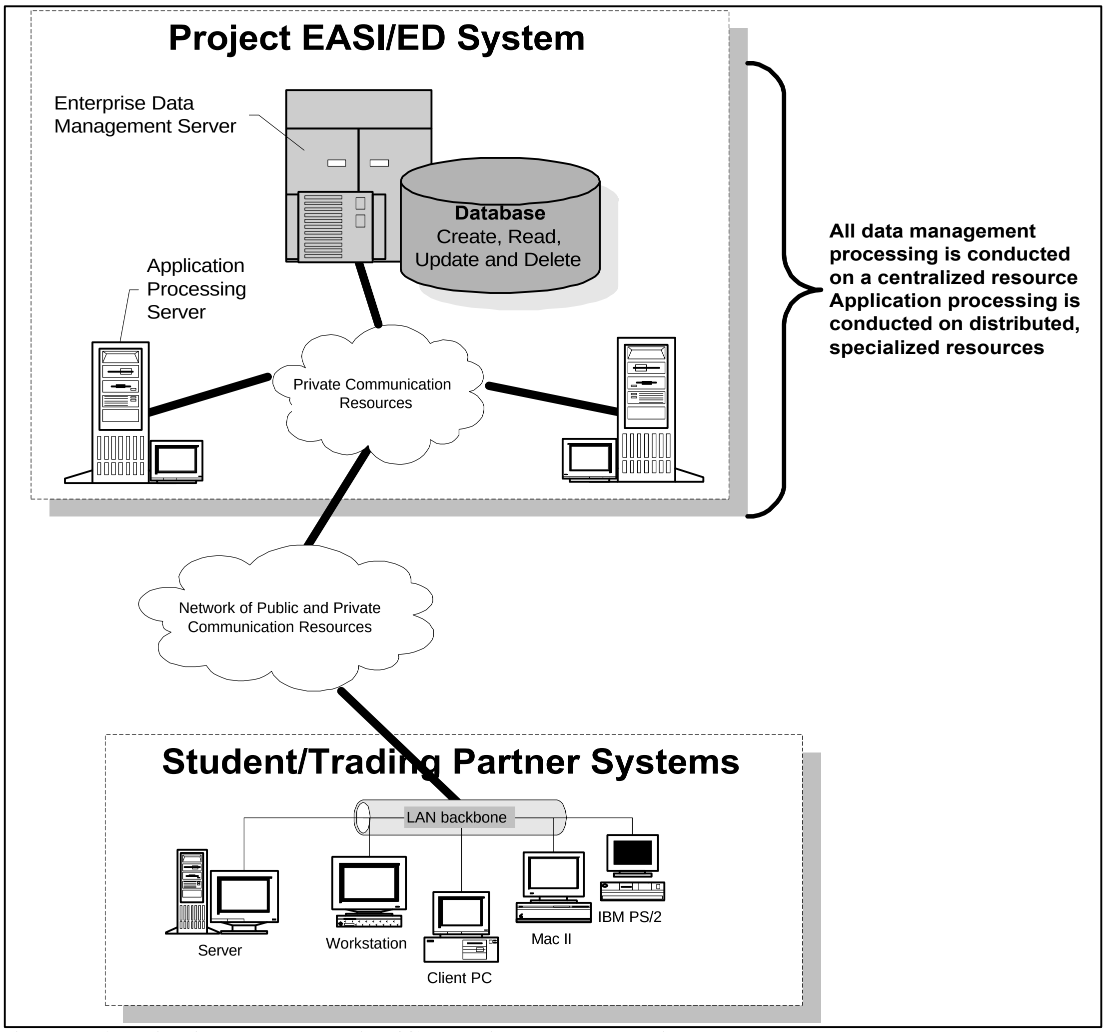

Figure 5-2. Distributed Processing/Centralized Data Architecture

5.1.3 Distributed Processing/Replicated Data for Consolidation

As illustrated in Figure 5-3, this architecture includes centralized decision support data management resources. The replication for data consolidation configuration facilitates collection of data from multiple primary sites, each of which support transaction processing activities. This “data consolidation” configuration is often useful in those situations where data may need to be regularly aggregated and reviewed, but distributed components need to be able to work without always being connected to the centralized site.

Figure 5-3 shows the Project EASI/ED processing distributed across multiple application processing servers. Updates to the database are made at multiple locations, with data “consolidated” into a central “read only” database. The system communicates with students and trading partner systems over a network of public and private communication resources.

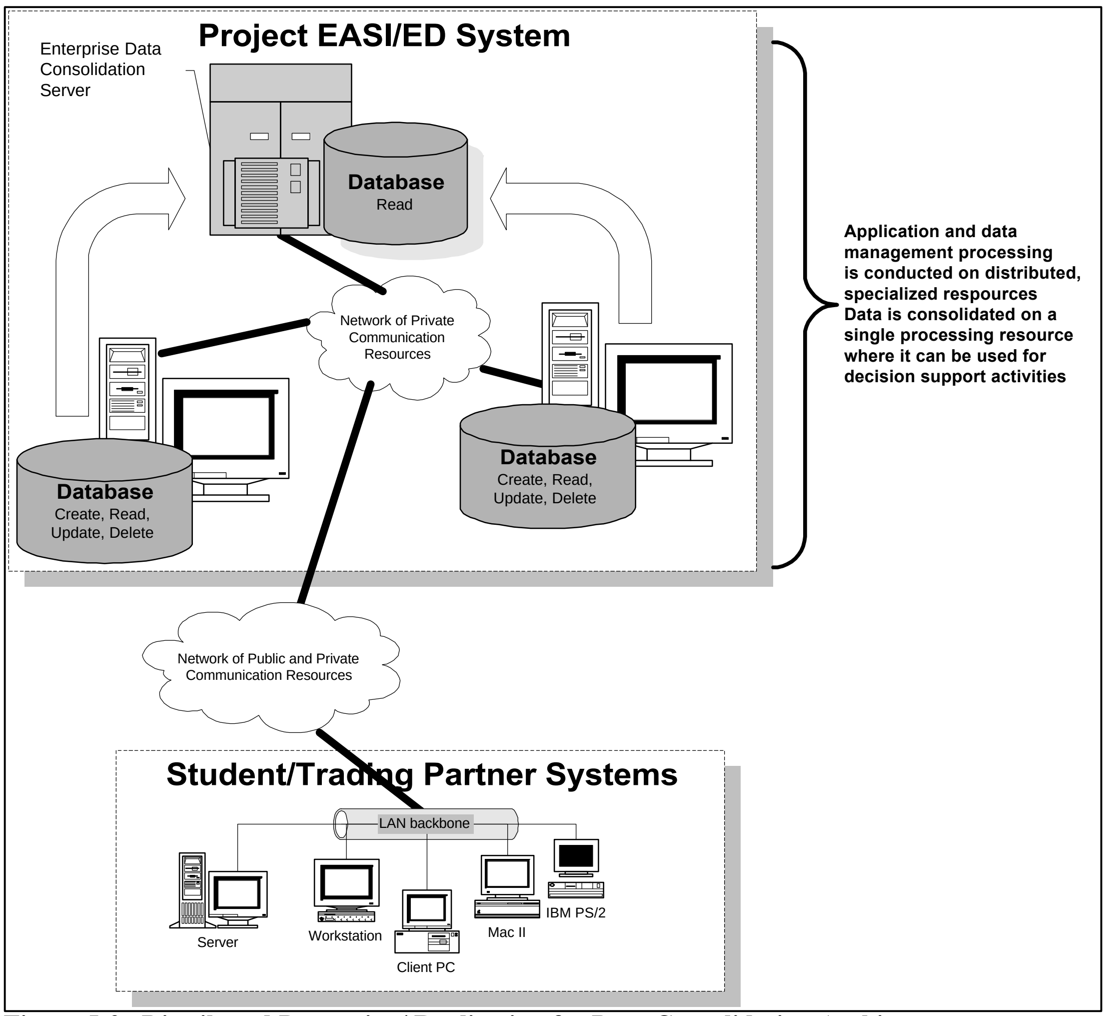

.**Figure 5-3. Distributed Processing/ Replication for Data Consolidation Architecture**

5.1.4 Distributed Processing/Replicated Data for Publication

As illustrated in Figure 5-4, this architecture includes centralized transaction processing data management resources. With primary-site replication for data publication, the primary site data resource copies data to multiple target data stores where it can be read. However, data is changed only at the primary site. The most simplistic example of this model is a single primary site that replicates all its data to a secondary system or to a set of identical secondary systems. In another, more complicated, configuration, portions of the primary site database could be copied to specified secondary sites, with each secondary site potentially receiving a different portion of the primary site database. This data replication configuration is often used to create hot-site backups and to populate distributed decision support system data repositories.

Figure 5-4 shows the Project EASI/ED processing distributed across multiple application processing servers. Updates to the database are only made on the central database, which copies data to a number of distributed “read-only” databases. The system communicates with students and trading partner systems over a network of public and private communication resources.

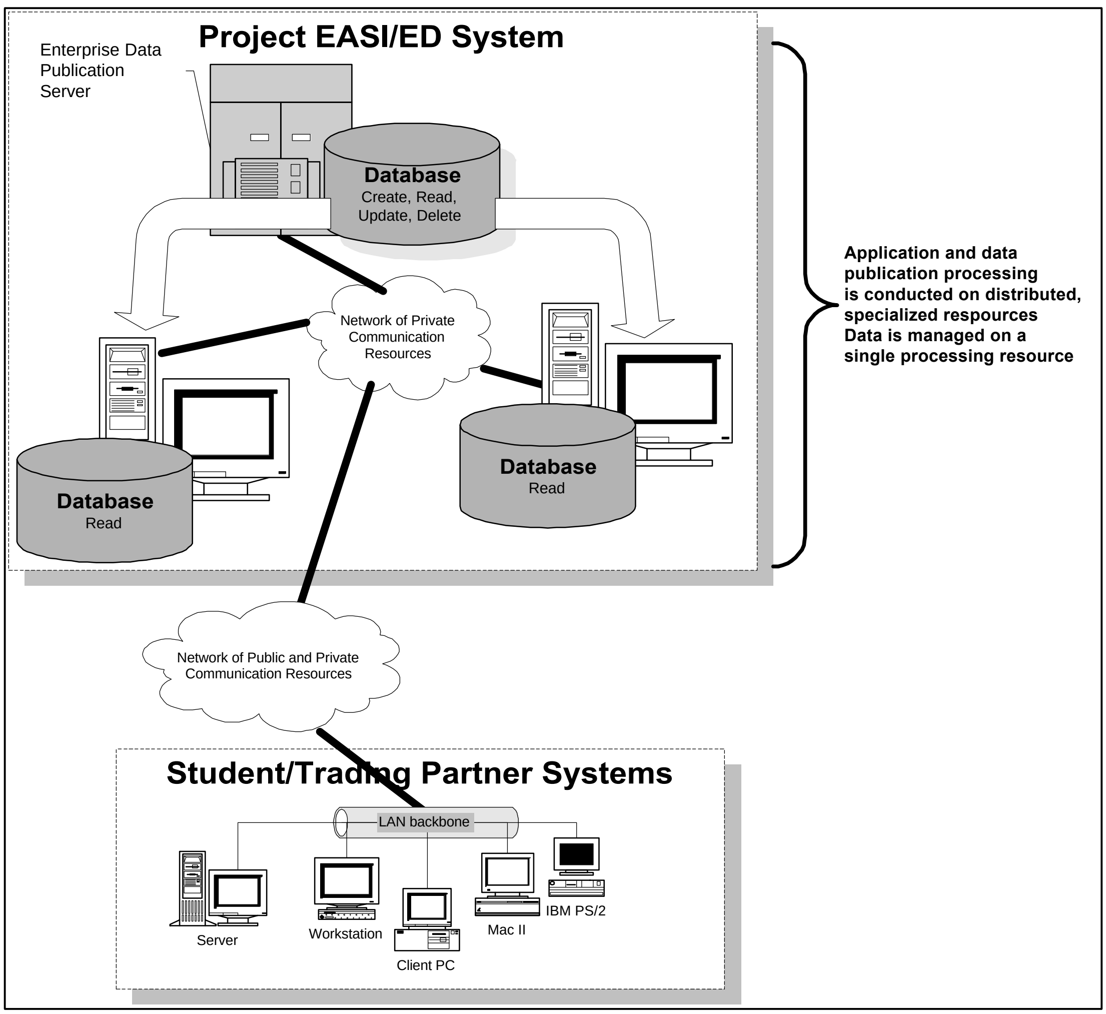

Figure 5-4. Distributed Processing/ Replication for Data Publication Architecture

5.2 Architecture Goals, Assumptions, and Constraints

This subsection identifies goals, assumptions, and constraints that affected this analysis. Goals, assumptions, and constraints are based on:

- Analysis of current Title IV systems, as described in Section 3.
- Evaluation criteria described in Section 4.
- Target system requirements, as defined within the *Project EASI/ED BARD* (July 1997).

5.2.1 Architecture Goals

The Project EASI/ED candidate architectures are designed to achieve the following goals:

1. Candidate framework architectures will provide a sound, integrated information technology framework for managing Project EASI/ED systems/subsystems and processes. They will provide a blueprint for developing and maintaining an integrated system capable of:

    - Ensuring that data collected and maintained within Project EASI/ED is accessible, understandable, and useful to ED and to ED's trading partners.
    - Delivering reliable, equitable, effective, and timely student financial aid and related services to students, prospective students, schools, lenders, guarantee agencies, and other enterprises associated with postsecondary education.
    - Streamlining, simplifying, and improving the accessibility of processes and data associated with postsecondary student financial aid delivery.
    - Reducing costs associated with managing and delivering postsecondary student financial aid.
    - Improving the postsecondary education community's ability to effectively reach and support a wider range of customers – thereby enabling organizations to establish and execute more successful business relationships with students.

2. Candidate framework architectures will enable Project EASI/ED to:

    - Improve delivery of student financial aid services.
    - Reduce annual system operations and maintenance budgets.
    - Provide ED, schools, guarantee agencies, students, prospective students, family members, and lenders with the accurate, complete, and timely information required to diminish fraud, waste, abuse, and mismanagement.

3. Candidate framework architectures will improve efficiency and reduce costs that result from:

    - Redundant and inconsistent data and data reporting.
    - Complex and expensive communications.
    - High administrative and operational costs.
    - Low-performance, high-cost technology.

5.2.2 Architecture Assumptions

Assumptions used for this analysis are:

1. Batch processing within Project EASI/ED will be similar to batch processing in the current Title IV systems. Candidate framework architectures include technologies that allow 65 percent of system processing to be performed in batch. (Processing modes are described in Subsection 2.2.2.)

This assumption is based on the average percentage of total processing currently performed in batch, as described in Subsection 3.2, Figure 3-2. Calculations associated with this assumption use system size, described as source lines of code, to determine the relative significance of current system batch processes. These calculations[1] are shown in Appendix G.

2. Candidate framework architectures will enable Project EASI/ED to reliably, and effectively support:

   - Management and disbursement of over $45 billion in financial assistance to over 8 million students annually.
   - Monitoring repayment of nearly $100 billion in Federal Direct and FFELP loans annually.
   - Processing 10 million FAFSAs annually to award Federal, state, and institutional aid.
   - Management and oversight of Title IV student aid assistance activities of more than 6500 lenders and 7,000 schools.

   This assumption is based on the General Accounting Office report titled: *Multiple, Nonintegrated Systems Hamper Management of Student Financial Aid Programs* (May 15, 1997).

3. Candidate framework architectures will enable Project EASI/ED to manage a total data volume in excess of 1.8 terabytes (1,764 GB). This assumption is based on the aggregation of data volumes managed by the current Title IV systems, as described in Subsection 3.3.3, Figure 3-5. This estimate is adjusted to reflect estimated current system data redundancy, as reported within the presentation titled *Future Title IV Delivery System Draft Systems Architecture and Plan of Action*. The Gunnison Consulting Group, Inc., presented this redundancy data to the Advisory Committee on Student Financial Assistance on March 14, 1996. This presentation reports that 13 percent of system master file data is redundant.

4. Candidate framework architectures will enable Project EASI/ED to support an average annual total data volume growth of 7 percent. This assumption is based on the average percentage of estimated annual data volume growth for Title IV systems, as described in Subsection 3.3.3 Figure 3-5. This calculation is shown in Appendix G.

5. Among other transactions, candidate framework architectures will enable Project EASI/ED to annually process the transaction volumes shown in Figure 5-5.

	Transaction Volume				
Transaction	1st Quarter (Jan-Mar)	2nd Quarter (May-June)	3rd Quarter (July-Sept)	4th Quarter (Oct-Dec)	Annual Total
Aid Applications	2,600,000	2,000,000	1,000,000	1,000,000	6,600,000
Aid Application Corrections	1,000,000	2,000,000	500,000	500,000	4,000,000
Aid Renewal Applications	3,000,000	2,000,000	500,000	500,000	6,000,000
Direct Loan Disbursements	1,296,059	1,927,601	896,319	1,005,362	5,125,341
Direct Loan Promissory Notes[2]	306,506	295133	1118352	781279	2,501,270
Aid Application Pre-Screening	10,000,000	3,000,000	3,000,000	10,000,000	26,000,000
Aid Grant Awards	1,500,000	800,000	1,200,000	1,500,000	5,000,000
Grant Award Confirmations	1,500,000	800,000	1,200,000	1,500,000	5,000,000
Total	21,202,565	12,822,734	9,414,671	16,786,641	60,226,611

Figure 5-5. Transaction volume of selected Title IV transactions

This assumption is based on data provided by ED in response to the Current Systems Questionnaire shown in Appendix E.

[1] These calculations do not consider disparities between programming language technologies.
[2] This data is based on data from the CDSI trending report dated November 1996. Specifically, this data represents Promissory Note transactions for the 01/96 to 12/96 time period.

6. Candidate framework architectures will enable Project EASI/ED to process all transactions according to the annual workload distribution shown in Figure 5-6.

	1st Quarter (Jan-Mar)	2nd Quarter (May-June)	3rd Quarter (July-Sept)	4th Quarter (Oct-Dec)
Transaction Workload Distribution	35.2%	21.29%	15.63%	27.87%

Figure 5-6. Quarterly workload distribution of Title IV transactions

This assumption is based on the workload distribution of those transactions described in assumption (5) and presumes that this transaction distribution of these transactions is representative of the transaction distribution for Project EASI/ED.

7. Candidate framework architectures will enable Project EASI/ED to support an average annual total transaction volume growth of 8 percent. This assumption is based on the average percentage of estimated annual transaction volume growth for the three most significant transactions processed by each of the current Title IV systems, as described in Subsection 3.3.3, Figure 3-4. This calculation is shown in Appendix G.

8. Candidate framework architectures will deliver total system productivity – defined as processed workload per unit of time (throughput) – equivalent to that provided by the current Title IV systems combined. Candidate framework architectures need to include technologies that, as a whole, are capable of processing over 1,179 MIPS (million instructions per second). This assumption is based on the total throughput capabilities of current Title IV systems, as described in Subsection 3.3.3, Figure 3-5.

9. Although 43 percent of the current Title IV systems are monolithic (as defined in Subsection 2.2.2.1), candidate framework architectures are not constrained by current system process distribution strategies. Rather, candidate framework architectures are based on target system requirements, as defined within the *Project EASI/ED BARD (July 1997)*.

10. Although 75 percent of the current Title IV systems are implemented, at least in part, in COBOL or a COBOL variant, candidate framework architectures are not constrained by current system technologies. Rather, candidate framework architectures, and technologies identified therein, are based on target system requirements, as defined within the *Project EASI/ED BARD* (*July 1997*).

11. Although 56 percent of the current Title IV systems are implemented, at least in part, in the MVS operating environment, candidate framework architectures are not constrained by current operating system technology selections. Rather, candidate framework architectures, and technologies identified therein, are based on target system requirements, as defined within the *Project EASI/ED BARD* (*July 1997*).

12. Candidate framework architectures must accommodate technically sophisticated trading partners, such as large universities and lending institutions, as well as schools, lenders, students, and other groups with little or no access to information technology. Therefore, candidate architectures support information sharing via Internet-based technologies, Electronic Data Interchange (EDI) facilities, electronic mail (E-mail) services, the US Postal Service mail (or equivalent), and telephony. This assumption is based on the General Accounting Office report titled: *Multiple, Nonintegrated Systems Hamper Management of Student Financial Aid Programs* (May 15, 1997).

13. Candidate framework architectures must enable users to access Project EASI/ED from a variety of disparate operating environments (e.g. Windows95, Macintosh, UNIX, and Windows NT.) Candidate framework architectures facilitate information sharing via commonly accepted, available, and affordable, standards-based technologies, protocols, and

data interchange formats. This assumption is based on the General Accounting Office report titled: *Multiple, Nonintegrated Systems Hamper Management of Student Financial Aid Programs* (May 15, 1997).

14. Candidate framework architectures must facilitate access to Project EASI/ED without restriction by user location, system access time, or specialized technical requirements. This assumption is based on the General Accounting Office report titled: *Multiple, Nonintegrated Systems Hamper Management of Student Financial Aid Programs* (May 15, 1997).

5.2.3 Architectural Constraints

Constraints that affected this analysis are:

1. Candidate framework architectures must manage data via relational database management technologies. This constraint is based on current Title IV system data management technology profiles, described within Subsection 3.3.3, which indicate that 100 percent of the current Title IV systems manage data, at least in part, with relational database management systems.

2. The Project EASI/ED networking infrastructure must be based on technologies that are capable of:

   - Leveraging many disparate local and wide area networking technologies to interconnect systems.
   - Operating with mature networking technologies, as well as with relatively new technologies, Fibre Distributed Data Interface (FDDI) or Asynchronous Transfer Mode (ATM).

   As a result, candidate framework architecture technologies will be networked via Transmission Control Protocol/Internet Protocol (TCP/IP). This constraint is based on current Title IV system network technology profiles, as described within Subsection 3.3.5, which indicate that 100 percent of the current Title IV systems communicate, at least in part, via TCP/IP.

3. To enhance system availability, fault tolerance, and performance, candidate framework architectures must include technologies that facilitate:

   - <u>Synchronous Communication,</u> which requires simultaneous availability of the client (service requesting) portion of the system and the server (service providing) portion of the system. Typically synchronous communication is "blocking." That is the service requestor must wait for the service response before processing can be continued. For example, synchronous communication might be used to request and provide student information that is needed for subsequent student aid administration processing.

   - <u>Asynchronous Communication,</u> which does not require the service provider to be available when a service request is issued. That is, application clients and servers do not have to be simultaneously available; rather, the services may be requested and then serviced at a later date or time. Asynchronous communication is typically not blocking. That is, the service requestor is free to conduct further processing while previously issued service requests are being satisfied. For example, Project EASI/ED might use asynchronous communication to request IRS validation of student income information. In this situation, the IRS batch cycles or other system workloads may not allow immediate request processing and response; however, because asynchronous and non-blocking communication is being used Project EASI/ED architecture is free to continue other operations while the IRS satisfies the request. Asynchronous non-blocking communication might also be used in situations where the service provider is not required to respond to the client at all. For example, a client may request a service provider to capture performance or audit information. In this situation, asynchronous communication would be appropriate, as (a) a response from the service provider is not required and (b) the service provider's ability to immediately satisfy requests should not affect client performance.

4. To reduce telecommunication costs and to improve accessibility, candidate framework architectures must leverage public network resources – the Internet – where possible, practical, and appropriate.

5. To support standardization, candidate framework architectures must require trading partners who electronically communicate data with Project EASI/ED via EDI to use appropriate X12 transaction sets including, but not limited to, those shown in Figure 5-7.

TRANSACTION SET ID	STANDARD TITLE	X12 REFERENCE NUMBER
130	Student Educational Record (Transcript)	X12.89
131	Student Educational Record (Transcript) Acknowledgment	X12.90
135	Student Loan Application	X12.198
139	Student Loan Guarantee Result	X12.265
144	Student Loan Transfer and Status Verification	X12.94
146	Request for Student Educational Record (Transcript)	X12.121
188	Educational Course Inventory	X12.322
189	Application for Admission to Educational Institutions	X12.321
190	Student Enrollment Verification	X12.264
191	Student Loan Pre-Claims and Claims	X12.276
194	Grant or Assistance Application	X12.372
198	Loan Verification Information	X12.359
810	Invoice	X12.2
820	Payment Order/Remittance Advice	X12.4

Figure 5-7. EDI X12 Transaction Set

6. Candidate framework architectures must require trading partners who access Project EASI/ED via Internet and World Wide Web technologies to use a standard Internet browser that supports Java™, JavaScript, Cookies, and the Secure Sockets Layer[3] (SSL) protocol.

7. Candidate framework architectures must require trading partners who access Project EASI/ED via X Window application components (X Clients) to do so using X Windows System Server (X Servers) technologies that:

- Comply fully with the X11R5 standard.
- Provide local X window manager services[4].
- Support the Motif Window Manager (MWM).
- Provide X client application startup over TCP/IP via the following methods/protocols: rexec, rsh, rlogin, and Telnet.

[3] SSL is a low-level encryption scheme used to encrypt transactions in higher level protocols, such as HyperText Transfer Protocol (HTTP), Network News Transfer Protocol (NNTP) and File Transfer Protocol (FTP).

[4] X Server technologies that provide local window manager services will allow X Window Manager processes to be executed using local (the user's) processing resources, thus reducing Project EASI/ED processing workloads.

- Provide X Display Manager Control Protocol services, including XDMCP-query, XDMCP-indirect, and XDMCP-broadcast.

8. Candidate framework architectures must include technologies that provide graphical human-computer interface services that are based on the X Window System and/or World Wide Web technologies.

9. Candidate framework architectures must require trading partners who electronically communicate with Project EASI/ED via electronic messaging technologies, such as E-mail, to use products that support the Simple Mail Transfer Protocol (SMTP) and the Multipurpose Internet Mail Extensions (MIME) protocol. SMTP is a server-to-server protocol that is widely used throughout the Internet to transfer E-mail messages between computers. The MIME protocol, which specifies uniform methods for handling data objects and rich text formatting within E-mail messages, is used to extend SMTP services. Because SMTP only supports the transfer of ASCII text and does not allow file attachments to be communicated, MIME is used in conjunction with SMTP to facilitate the communication of messages containing voice, graphics, video, and EDI file attachments.

10. Candidate framework architectures must include E-Mail server technologies that support the Interactive Mail Access Protocol Version 4 (IMAP4). IMAP4 is a client-to-server protocol used to transfer descriptive mail message headers (along with other message bodies and attachments that are explicitly requested) between the mail hub and electronic mail clients (user agents [UAs]). IMAP4 allows E-Mail clients to operate when disconnected from mail servers. IMAP4 also reduces network overhead associated with transferring entire mail files, which often include unwanted content, between the mail hub and the UA. Using IMAP4 users can review, select, and delete mail file contents without having to first download the entire mail file from the mail server.

11. Candidate framework architectures must include electronic mail server technologies that support the Post Office Protocol (POP). POP is a standardized batch protocol that is almost universally used to collect mail from Internet Service Providers (ISPs). POP is also widely used to move mail from mail hubs to the systems providing UA services – electronic mail clients.

12. Candidate framework architectures must include relational database management system (RDBMS) and transaction processing (TP) technologies that support X/Open-compliant transaction management. In particular, RDBMS and TP technologies will coordinate and communicate transactional information via the X/Open defined XA interface. The XA interface is a standardized interface to manage transactions via the two-phase commit protocol. RDBMS and TP technology XA interface support will facilitate the distributed transaction management involving multiple, heterogeneous databases.

13. Candidate framework architectures must include RDBMS technologies that are ISO SQL-92 (also referred to as SQL2) compliant. SQL-92 was published by the International Standards Organization and ratified in 1992. The standard is a superset of the SQL-89 standard, which was originally published in 1986 and amended to include Embedded SQL in 1989. SQL-89 defines a standardized SQL Data Definition Language (DDL), a Data Manipulation Language (DML), referential integrity, and embedded SQL mechanisms. In addition to these services, SQL-92 compliant technologies support additional features, including:

    - SQL client/server connections and sessions.
    - Embedded SQL support for the C, Ada, and other programming languages.
    - Binary large object (BLOBs), timestamp, varchar and other new data types.
    - Dynamic SQL, which may be generated at system runtime.

5.3 Architecture Strategies

The purpose of the Project EASI/ED architecture analysis is to establish agreement among the involved organizations regarding the basic (framework) process and data distribution model for the system. With this in mind, the team populated each of the candidate framework architecture models (described in Subsection 5.1) with specific technology that facilitates a complete evaluation. A consistent suite of technologies was used across all four candidate frameworks to ensure that the evaluation would focus on the relative merits of the distribution models, rather than on the merits of various vendors' products. As was stated previously, the specific technologies and configurations presented in this report are intended to support the framework architecture evaluation only. While they may be used as a basis for moving forward with the later stages of architectural analysis, Project EASI/ED is not limited to these architectures at this point.

This subsection presents the overall strategies used to select specific technologies (e.g., hardware, software, telecommunications) used to populate the candidate frameworks. These strategies form the jumping off point for defining and understanding the specific sets of technology described in Appendix C, which in turn are used in the detailed descriptions of the candidate framework architectures presented in Subsection 5.4.

5.3.1 Common Operating Environment

Candidate framework architectures are based on open systems technologies as implemented by Sun Microsystems. All operating environments within the architectures will be Sun Ultra class machines running the Solaris[5] operating system. These technologies minimize integration complexity, enhance security, simplify system management, and reduce training requirements. Candidate framework architectures will use homogeneous (common) operating environment technologies.

The Sun Ultra class machine strategy was selected based upon the following rationale:

- Sun Microsystems holds a leading position in the enterprise and Web server markets, and more than 12,000 applications are available for Ultra class machines.

- Sun's Ultra High Performance Computer (HPC) servers provide highly scaleable architecture. Product offerings in the Ultra server family range from a powerful workgroup server (the HPC 2) to a high-end enterprise computer center server (Starfire Enterprise HPC 10000) that is designed to offer the level of service traditionally provided by mainframe technology. Sun Ultra servers can be clustered, providing linear scalability in excess of that offered by uniprocessor or Symmetric Multi-Processor (SMP) solutions, but with fewer of the systems management complexities inherent in Massively Parallel Processing (MPP) products.

- The Sun Ultra series provides a wide range of High Availability (HA) server solutions (see Subsections 5.3.5) that facilitate advanced data, processing, applications, and network fault tolerance. High availability on-line transaction processing services for BEA Tuxedo, the most widely used open systems OLTP monitor, are available exclusively for the Sun Microsystems Ultra Enterprises server class machines. Sun servers typically achieve availability in the 99 percent range. The use of clustering technologies can improve availability into the 99.98 percent range (equivalent to less than 2 hours of downtime per year).

[5] Solaris was voted “Best Internet Platform” by *Network Computing* in February, 1997. Solaris was named “Best UNIX Operating System” by *Network VAR* in January, 1997.

The Solaris-based common operating environment strategy was selected based upon the following rationale:

- A common operating system minimizes security and security management complexity common in heterogeneous and distributed environments. A common operating system allows for use of common user profiles and access control lists across multiple systems. This simplifies administration and facilitates single sign-on security solutions.

- More than 12,000 applications are available for the Solaris operating system on SPARC. Furthermore, Solaris is considered a “first port” for most popular open systems transaction processing and management information system technologies.

- Sun provides a comprehensive set of:

  - High availability server solutions (see Subsections 5.3.5) that facilitate advanced data, processing, application, and network fault tolerance.

  - Integrated system management, directory, and security services, which facilitate centralized and simplified distributed systems management. Examples include the Solstice software suite, which includes AutoClient (see Subsection 5.3.7), DiskSuite (see Subsection 5.3.5), and Internet Mail Server.

  - Network File System solutions, which facilitate dynamic, server-based software deployment and configuration management, as described in Subsection 5.3.7.

  - The Solaris technology suite is widely used and is recognized for reliability, availability, and scalability.

5.3.2 Application Security

Candidate framework architectures “expose” Project EASI/ED to the Internet. To accommodate the Internet, Project EASI/ED requires multilevel security solutions (MSS) to safely manage information with different sensitivities, while preventing unauthorized users from access to information. Candidate framework architectures facilitate convenient identification, authentication, access control, authorization, confidentiality, integrity and non-repudiation.

To satisfy these needs, candidate framework architectures use application specific security mechanisms. Application coupled security is usually provided by applications with specific security requirements and is typically implemented as a transport layer technology[6] (such as Netscape’s SSL). EDI, E-mail, and Web-based processing will provide application-specific security services. Specific product implementations of these services are detailed in Appendix C – Candidate Architecture Technologies.

---

[6] The transport layer is defined within the Open Standards Organization’s Open System Interconnection (OSI) model. For more information about the OSI model refer to the TCP/IP abstract within Appendix B.

5.3.3 Secure Virtual Private Networks

As of 1996, the Internet connected more than 20 million users in 70 countries and was growing at an estimated 10 to 20 percent per month[7]. For high-priority corporate network backbones[8], Internet-based Wide Area Network (WAN) solutions can reduce those costs typically associated with dedicated, leased line solutions,, by nearly 50 percent. Similarly, low-priority branch office use[9] of Internet-based WAN technologies can reduce cost by as much as 23 percent. Given this, to reduce telecommunication costs and to improve accessibility through Project EASI/ED, candidate framework architectures leverage the Internet where possible, practical, and appropriate. Candidate framework architectures leverage application security solutions and public telecommunication resources to logically implement secure "virtual private networks" for EDI, E-mail and Web-based processing.

5.3.4 Data Fault Tolerance

Candidate framework architectures include reliable storage subsystem technologies that are based on Redundant Array of Independent Disks (RAID) technology. Six RAID strategies can be used to reduce costs, improve fault tolerance, and/or enhance performance. Candidate framework architectures use RAID 5 for applications dominated by read access operations, such as data warehousing, decision support, and other analytical processing. RAID 1+0 configurations are used in applications requiring enhanced fault tolerance and involving significant random and sequential write operations, such as on-line transaction processing.

- RAID 5 strategies use disk striping to distribute disk I/O across multiple disks and disk controllers. Additionally, RAID 5 configurations provide efficient fault tolerance by calculating and storing parity data, which can used to recover data in the case of failure, across multiple disks.

- RAID 1+0 (also referred to as RAID 10) strategies use disk striping to distribute disk I/O across multiple disks and disk controllers. Additionally, RAID 1+0 configurations provide complete data redundancy through disk mirroring. This ensures superior fault tolerance and recoverability.

These RAID configurations facilitate data resource "hot sparing" to allow components to be continuously available. Availability is facilitated through the use of redundant system components configured to automatically replace failed primary components – thereby mitigating system interruption. Fault tolerant RAID configurations allow mirrored data, for example, to be automatically substituted for primary data resources. Additionally, candidate framework architectures use storage array subsystems capable of "hot swapping," disk resources. Hot swapping allows a failed component to be replaced.

---

[7] As reported within US Computer's report titled: *Internet-Based Secure Virtual Private Networks: The Cost of Ownership* (April 1996)
[8] High-priority backbone networks are mission-critical networks supporting key corporate applications. These networks are designed to have consistent, high levels of bandwidth, and low latency. Typical applications for such networks are inventory, purchasing, and general accounting.
[9] Low-priority branch office networks are networks where reliability requirements are not as extreme, such as branch office networks primarily used for E-mail, file transfers, or low-priority database access.

5.3.5 High Availability Processing

Candidate framework architectures include highly available application processing facilities. Highly available environments provide fully redundant processing resources, so recovery from failure takes only seconds or minutes.

Candidate framework architectures employ server "clustering" technology to provide redundant paths among all systems, to all disk subsystems, and to all external networks. As a result, single points of failure are eliminated and services are highly available. Figure 5-8 illustrates a server cluster.

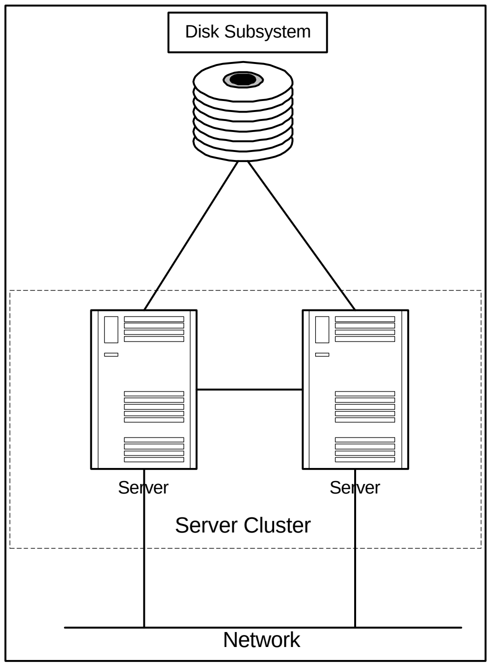

Figure 5-8. Server Clustering

To implement clustered solutions, configurations include:

- Two or more redundant server systems (nodes).
- Fault tolerant storage systems.
- Redundant network interfaces.
- Management software for unified control of cluster components.
- Hardware and application monitoring, and failure detection, and failover ("hot sparing") management software.

Candidate framework architectures use Sun Microsystems's clustered Ultra Enterprise server configurations to facilitate the following highly available system services:

- E-mail.
- EDI.
- On-line transaction oriented relational database management.
- Network file systems.

These system services are made highly available through the following technologies:

- *Ultra Enterprise Cluster HA Servers*, which are designed for environments requiring high availability data, file, and application services. These servers ensure rapid detection and recovery from any hardware, network, operating system, or application software failure.

- *Ultra Enterprise Cluster PDB (Parallel Database) Servers*, which are designed to support parallel database operations, such as those offered by “cluster-aware” RDBMS technologies like the Oracle 8 Parallel Server. These servers facilitate enhanced performance via multi-node parallel processing, as well as rapid detection and recovery from any hardware, network, operating system, or application software failure.

- *SPARCstorage* arrays, which provide fault tolerant data storage services, as described in Subsection 5.3.4.

- *Solstice DiskSuite 4.0*, which is a “logging” data volume management software solution that ensures data integrity on all file systems. This technology also allows both nodes in the cluster to share available storage resources – allowing the surviving node in a cluster to assume the identity (via IP failover and other services) of the failed member and transparently service data requests. Additionally, DiskSuite allows read/write operations directed at a failed disk to be automatically and transparently redirected and serviced by mirrored storage.

- *Solstice HA 1.3*, which provides the intelligence to detect and recover from a variety of hardware, operating system, application, and network failures. Specifically, HA 1.3 uses “fault probes” to monitor and detect failures, including those associated with NFS and Oracle.

- *HA Extensions for BEA Tuxedo*, which provides for automatic and transparent resubmission of “in-flight” transactions, which have been begun, but not committed, when a hardware, network, or database service fails. *HA Extensions* monitor the availability of database services and in cases of failure will rollback uncommitted transactions and resubmit these transactions to the surviving database server cluster node (without the users’ knowledge or intervention).

5.3.6 Network Media

Ethernet is the most popular data link protocol[10] in use today and is the most common way to implement Local Area Networks (LANs). This popularity is largely the result of the technology’s ability to provide inexpensive, flexible, high speed, and maintainable data link services via Carrier Sense Multiple Access with Collision Detection (CSMA\CD) technology.

With CSMA\CD, any station is allowed to transmit data at any time, as long as the network is not occupied by transmissions from other stations. When two or more stations transmit data simultaneously, a collision occurs. When a collision occurs, the stations involved wait a random period of time before attempting to transmit again, after first listening to make sure the network is available. However, because all stations on a Ethernet network share the available bandwidth[11], as the network traffic increases, so do the number of network collisions – often resulting in lower overall throughput[12]. To mitigate performance problems associated with this phenomenon, the Candidate framework architectures use Switched, Fast Ethernet.

---

[10] For more information regarding data link protocols refer to the Network Infrastructure abstract in Appendix B.
[11] Bandwidth describes the amount of data that can be transmitted on a physical medium at one time.
[12] Throughput is the measure of data that is transmitted between to points. Throughput is a function of bandwidth and data transmission rate.

Using switch technology, Switched Ethernet provides the network's entire bandwidth to each station for small, alternating intervals of time. Using this technique, stations on the network do not share the bandwidth, as they do with Shared Ethernet.

With the arrival of economical, powerful desktop technology and the subsequent distribution of complex applications, organizations realized a need for increased network bandwidth. With data rates of up to 100Mbps, Fast Ethernet satisfies these needs. This enables Fast Ethernet to provide data rates required by Project EASI/ED without the complexity or expense associated with alternative technologies, such as FDDI technology.

5.3.7 Centralized Software Management and Dynamic Deployment

The client platforms within many Web-based architectures are very "thin" – hosting only a Web browser and no custom code. As a result, system modifications and configuration changes are only made to system components residing on Web, application, and data management servers. Likewise, most new system "features" need only be deployed to a very small number of machines – the servers – where they can be immediately accessed by system users. All clients access applications via the system's Web server and there is no need to worry about which clients are using which version of an application.

To realize the benefits of centralized software management and dynamic deployment - not just with Web-based components, but throughout the entire system – candidate framework architectures employ Network File System (NFS) -based technologies. NFS technologies simplify administration, increase availability, and centralize management operations. Where possible candidate framework architectures use Sun Microsystems's Solstice AutoClient 2.1 – a Cache File System (CacheFS)-based NFS technology – to implement "swappable" servers.

NFS is a set of machine- and operating system-independent file transfer protocols that support file transfer over heterogeneous networks[13]. NFS allows portions of one system's file hierarchy to be transparently mounted within another system's hierarchy. This allows operating environments to transparently share file systems in multi-vendor networks and operating environments.

CacheFS enhances NFS by providing a general purpose file system caching mechanism that improves NFS server performance and scalability by reducing server and network load. Where NFS only provides clients with "pointers" to shared server-based files and files systems, CacheFS provides the capability to store a local copy of any file that is requested from an NFS server. In doing so, CacheFS eliminates latencies associated with retrieving files across a network, thus freeing the server to support other activities.

Through Solstice AutoClient 2.1 users see what they would see working on a standalone system. At system boot, the AutoClient software loads the operating system, network-mounted applications, and data from the server through NFS mounts. All these are stored in a local cache and memory is swapped to the local disk for speed. Any locally cached data that is modified is recorded in the cache and then (if configured to do so) written through to the server, where the corresponding back file is updated. If the user accesses executables, applications, or data that are not locally cached, these are downloaded to the client's file system on demand.

Using AutoClient technology, candidate framework architectures mount centrally located application software from a central location – the NFS server – onto multiple application servers. This provides a number of benefits.

- Eliminates need to back up clients as non-transient data is maintained on the NFS server.
- Distributes versions of application software to servers through NFS mounts, thus reducing administration costs.

---

[13] NFS uses a technique called Remote Procedure Call (RPC) to provide operating system-independent network functionality. RPCs are discussed in the Appendix B middleware abstract.

- Allows replacement of faulty servers and return on-line in minutes, as servers are automatically configured through the AutoClient boot process.
- Allows central software maintenance and administration.

Consider the practical benefits of this strategy. As Figure 5-9 illustrates, the NFS server allows application servers (NFS clients) to dynamically load software that is managed, administered, and maintained at a centralized location – the NFS server. This allows not only system components to be “swapped” out in cases of failure, but also allows entire application servers to be quickly replaced. Furthermore, consider the vulnerability of publicly accessible systems such as Web servers. As the previous illustration depicts, software running on these servers can be dynamically loaded through the security firewall. This not only allows application software to be maintained within a secure environment, but in cases of security “breach” on the Web server, also allows the server to be recycled and dynamically reconfigured – quickly correcting unwarranted system modifications.

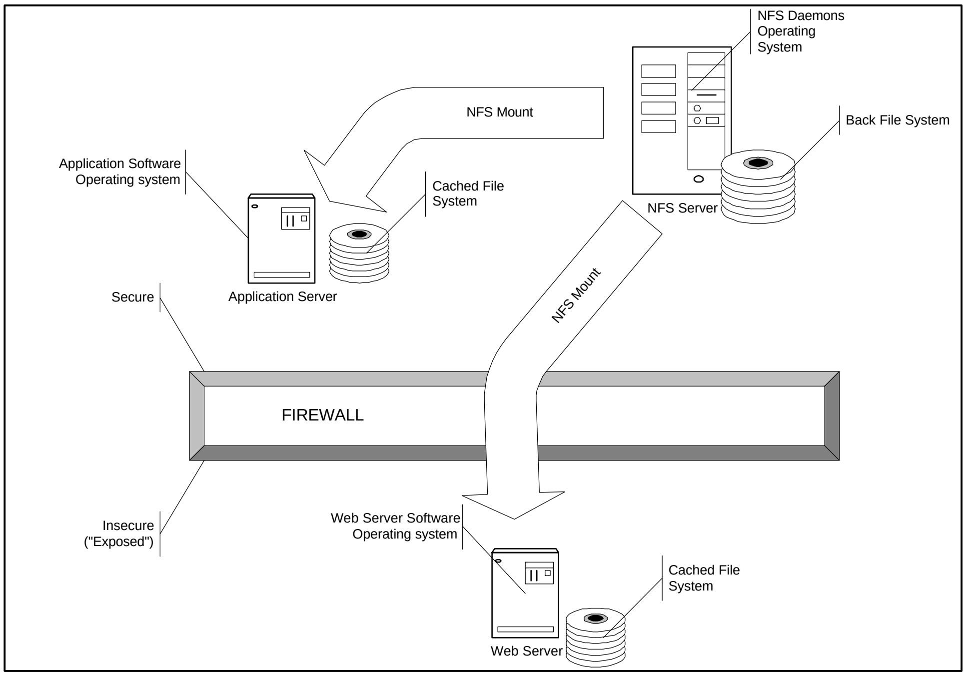

Figure 5-9. Mounting NFS on Application and Web Server

5.3.8 Centralized Execution and Remote Display of Application Clients

Candidate framework architectures employ technologies that minimize the amount of system application logic distributed to client platforms. By minimizing client platform application distribution, configuration changes can be implemented by modifying application components that reside on Project EASI/ED servers. This significantly reduces complexities associated with deploying new software to all system users.

This strategy is achieved largely through the use of Web-based technologies, as described in Subsection 5.3.7. Candidate framework architectures will employ X Window System technologies and centralized “application display” servers, as illustrated in Figure 5-10.

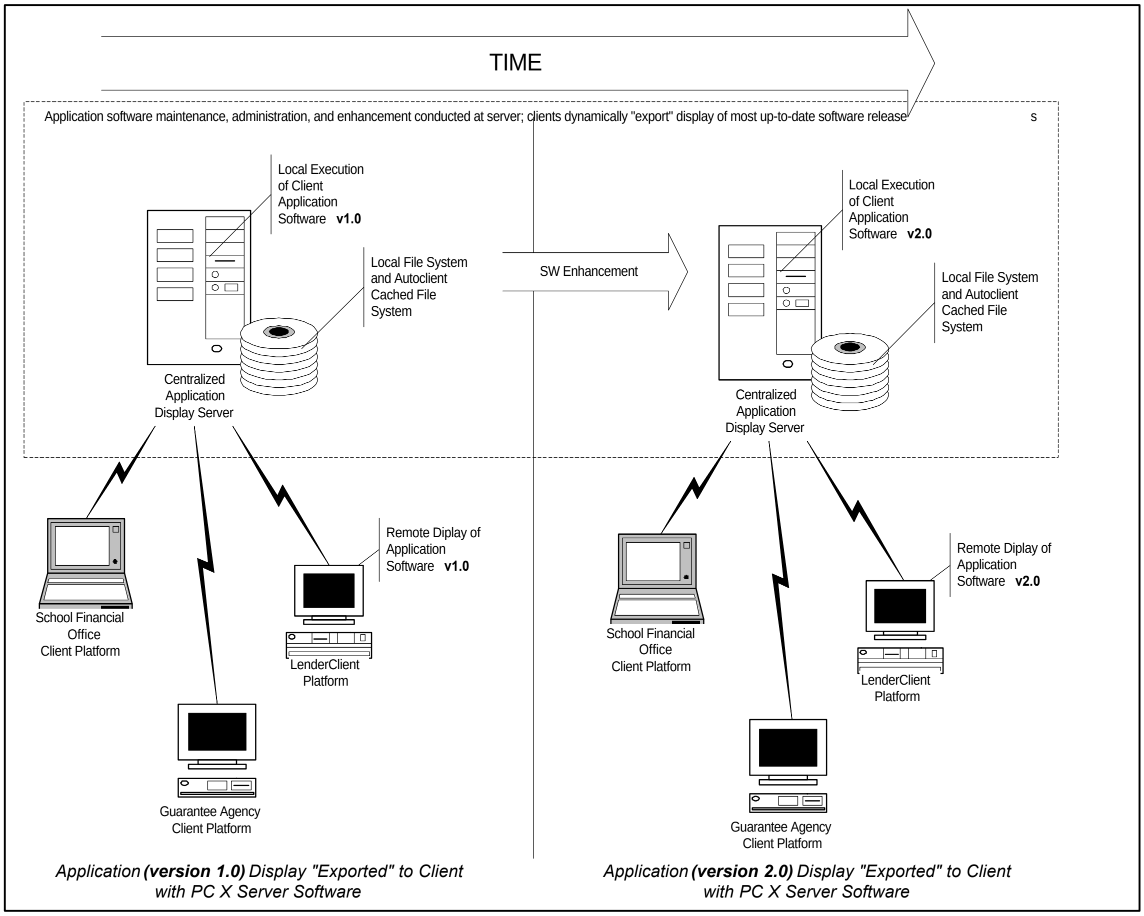

Figure 5-10. Centralized software distribution and maintenance strategy

These technologies allow client application software to be executed from a highly available enterprise class server environment, while facilitating the remote display of the application software's graphical user interface. This strategy greatly reduces system administration complexity and cost, while simultaneously providing remote users with access to the most "up-to-date" application software.

5.3.9 Network Firewall

Candidate framework architectures use firewalls to provide Project EASI/ED with full, transparent, Internet connectivity using the entire range of Internet protocols, while also ensuring network security. To provide network security, the firewall must be the only point through which Project EASI/ED and the Internet can connect. Firewalls implement an organization's network security policy by inspecting all data packets entering (and leaving, if desired) the network from the Internet, and deciding whether or not the packets should be permitted to pass through.

In general there are two types of firewall: network and application. (However, the latest firewall technologies are blurring the distinction between the two.) Network-level firewalls make their decisions based on network address (the originating host and port number, along with the destination host and port number). A simple router would be such a network-level firewall, as it cannot make complex decisions based on knowledge of an application's networking requirements. Application-level firewalls generally act as intermediaries or proxies for traffic between internal and external networks; no traffic is permitted directly between these networks. Application proxies make their decisions based on knowledge about specific application traffic, which means that detailed audit information can be provided. Many application-level firewalls act as an obvious gateway between internal and external networks, and although this may require some training for users, this can also be used to remind them that they are passing from one zone to another. Newer firewalls often make access to the external network fully transparent to users inside.

The firewall will protect any number of HTTP servers behind the gateway through authentication and address translation. The administrator determines how each user is authenticated and what resources are available.

A firewall can conceal the network's internal IP addresses from the Internet and translate internal addresses that are invalid internet addresses, in both cases by assigning an IP address from a pool of addresses set aside for the purpose of being the network's interface to the Internet. Address translation may also be used to implement "one-way-routing," so that there is no route from the outside to an internal network or host.

5.4 CANDIDATE FRAMEWORK ARCHITECTURES

This subsection defines in detail the four candidate framework architectures evaluated in this report.

Subsection 5.4.1 **-** Centralized Process/Centralized Data.
Subsection 5.4.2 **-** Distributed Process/Centralized Data.
Subsection 5.4.3 **-** Distributed Process/Replicated Data for Consolidation.
Subsection 5.4.4 **-** Distributed Process/Replicated Data for Publication.

Each subsection comprises the following elements:

- A brief overview of the major features of the candidate framework architecture
- A schematic diagram showing the hardware components comprising the architecture and their interconnection
- A list of the hardware and software components comprising the architecture. These components are described in more detail in Appendix C.
- A calculation of the storage capacity required for the architecture

The specific products with which each architecture is populated appear in the *Project EASI/ED TVTA Report* for cost comparison purposes only, and the inclusion of a particular product does not constitute a recommendation that that product be used to implement Project EASI/ED. A full vendor evaluation will be conducted in a later phase of Project EASI/ED to recommend specific products.

5.4.1 Centralized Process/Centralized Data Architecture

This architecture includes centralized data and application resources, which facilitate the execution of all data management, application, and presentation processing from a single computing resource. This architecture is typically characterized as:

- **Secure -** System processing resources can be physically secured in at a single location and logical security of the system is not complicated by the need to integrate multiple, heterogeneous security solutions that might be required without a single, centralized operating environment.
- **Performance constrained -** The use of common processing resources to simultaneously support DSS and transaction processing is often prohibitive.
- **Vulnerable -** All system operations can be catastrophically affected by single resource failure.
- **Fault intolerant -** Synchronous availability of system components mandated by single, monolithic operating environment.
- **Growth constrained -** System size, throughput, and workload capacity limited by scalability of centralized processing resource.
- **Manageable -** Software deployment, configuration management, maintenance, and performance monitoring can be managed from a single location without the complication of remote systems administration.
- **Inflexible -** System growth and service support constrained by single hardware environment and limited vendor or third-party product offerings.

With these characteristics in mind, a fully centralized framework architecture for Project EASI/ED was designed to mitigate weaknesses and leverage strengths intrinsic to centralized process/centralized data architectures. This architecture is detailed in Figure 5-11.

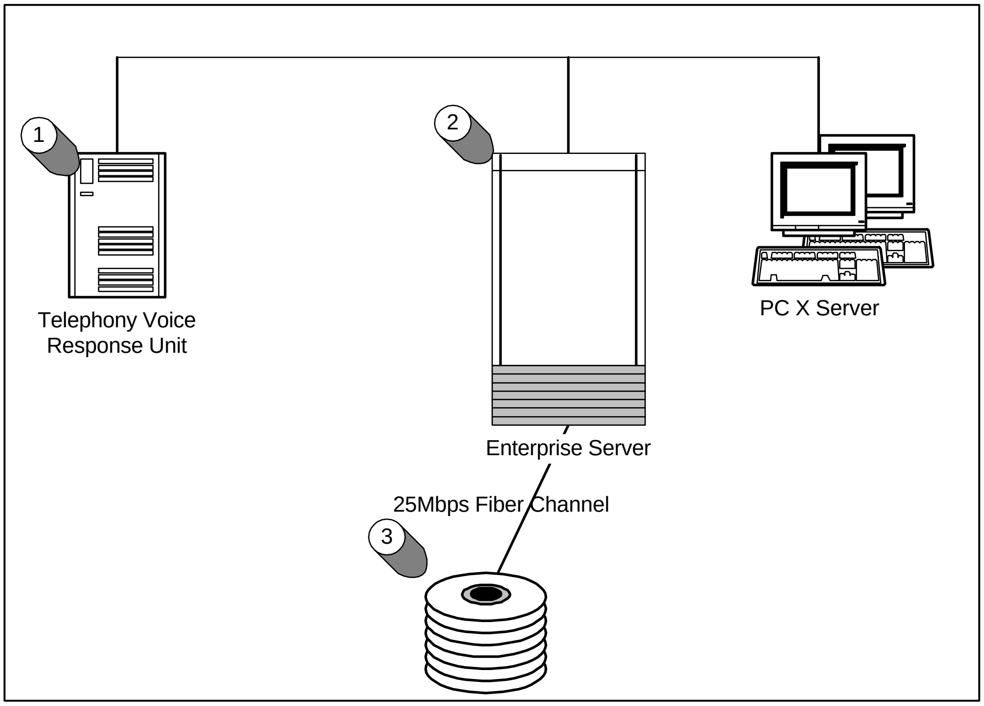

Figure 5-11. Centralized Processing/Centralized Data Architecture

The previous illustration describes interfaces between hardware components and complements the software interface descriptions provided in Appendix C.

Unlike other candidate architectures described in this subsection, this architecture, does not allocate software components to “specialized” hardware resources. Rather, a single computing resource – a Sun Ultra Enterprise 10000 server (UE 10000) – is used to support all system operations. With the exception of processing performed by the VTK 3000 Voice Response Unit, which provides interactive voice response services, all system processing, reporting, data management, user interaction, information sharing, and other operations are processed by the UE 10000.

As is the case with all of the candidates, this architecture allows users to invoke a common set of system services, such as the Tuxedo OLTP monitor, via telephony, Web-based, and LAN/WAN X Window System-based application software. However, unlike other candidate architectures, all of these system services will be served from a single operating environment – the UE10000.

In addition to centralized processing, this architecture also centralizes data management activities. A single data repository is used to simultaneously support both transaction management and decision support operations. As a result, this architecture does not provide data warehouse technologies or services. Rather, data is managed via Oracle’s RDBMS technology and will be “served” from the UE 10000.

The UE 10000 is configured to have 100 percent hardware redundancy. That is, the server configuration used within the centralized process/centralized data architecture includes redundant control boards, system boards, disk storage, power and cooling subsystems, peripheral controllers, etc. Additionally, fault tolerance system services will be implemented via the UE 10000’s Dynamic System Domains (DSD) technologies.

In many data centers, mission-critical applications are isolated from other possibly “offensive” system operations by establishing specialized servers for certain applications. In a mainframe environment, system partitions may be established for the same purposes. The UE 10000’s DSD provides similar services – enabling multiple systems to be created within a system via logical partitions (domains).

With DSD, each domain is a fully functional, self-contained server. Domains comprise one or more system boards. Each domain has its own Central Processing Unit (CPU), memory, I/O and network resources, and runs its own instance of the Solaris operating system. Because each domain has its own resources, it is, to a large degree, physically separate from other domains. As a result, all software errors originating within a domain are isolated. This isolation is largely facilitated through the use of completely independent instances of the Solaris operating system. Additionally, most hardware errors originating within a domain will not affect any other domain. For example, if a processor fails in a domain, no other domain will be affected. However, certain components, such as the system control board, are shared system-wide and failures with these “shared” components do affect the entire system and all domains.

The following table, Figure 5.12, describes the hardware and software components that comprise the candidate centralized process/centralized data architecture.

Reference	Vendor/Product	Description	Quantity	Hardware Specification
1	VoiceTek, Inc[14] VTK 3000	Interactive Voice Response Unit Hardware Hardware Server	3	- 1 Pentium processor - 166Mhz clock speed - 32MB RAM - 4 GB of Disk - SCO UNIX 5.0.2 - T1 trunk interface - 100BaseT LAN interface
	VoiceTek, Inc. Generations Telephony Server Platform v4.0	Telephony – Interactive Voice Response Unit Software	3	
2	Sun Microsystems Ultra Enterprise 10000	Highly Available, Enterprise Class Server	1	64 UltraSPARC II processors running at 250 Mhz clock speed with 64 GB RAM running Solaris 2.6
	Solstice HA v1.3	High Availability Hardware Monitoring and Fault Detection Software	1	
	Solstice PDB v1.2	High Availability, Parallel Processing Hardware Monitoring and Fault Detection Software	1	
	Solstice: DiskSuite v4.0	High Availability Volume Management Software	1	
	BEA Systems, Inc. HA Extensions for BEA Tuxedo	High Availability, Hardware Monitoring and Fault Detection Software for Tuxedo	1	
	VoiceTek, Inc. Generations Runtime Server Platform v4.0	Telephony – Runtime Application Server	3	

Figure 5-12. Centralized Process/Centralized Data Architecture Components

[14] VoiceTek, Inc. reports that several organizations, including Atena, Nynex Telesector Resources Group, and Oxford Health Plans, are successfully using similarly configured VTK 3000 voice response units to deliver telephony services. Each VTK 3000 Telephony server is networked with a digital PBX facility via a T1 connection and will simultaneously support 120 calls. As a result the recommended configuration will support 360 simultaneous calls.

Reference	Vendor/Product	Description	Quantity	Hardware Specification
2	Premenos Corporation Templar Enterprise	EDI Security and SMTP/MIME Delivery Software	1	
	Premenos Corporation EDI/Open	EDI X12 Translation Software	1	
	Oracle Corporation Oracle 8 Parallel Server	Relational Database Management System Software	1	
	Oracle Corporation Oracle 8 Server	Relational Database Management System Software	1	
	Netscape Communication Corporation Enterprise Server v3.0	Web Server Software	1	
	Netscape Communication Corporation Directory Server[15]	Web Directory Server Software	1	
	BEA Systems, Inc. BEA Tuxedo v6.x	Online Transaction Processing Monitor	1	
	BEA Systems, Inc. HA Extensions – HA Client Wrapper	High Availability, In-flight transaction recovery and resubmission software	1	
	BEA Systems, Inc. BEA Jolt v1.1	Web Gateway for Online Transaction Processing Monitor Services	1	
	SAS Institute, Inc SAS	Statistical and Quantitative Analysis Software	1 (site license)	
	Sun Microsystems, Inc. Solstice Internet Mail Server v2.0	Internet Electronic Mail Server Software	1	
	DNS	Domain Name Server Software	2	
	DAZEL Corporation DAZEL Output Server[16]	Output Management System Software	1	
3	SPARCstorage Array 214 RSM[17]	Data Center Class Storage Array Subsystem	36	- 1x110 Mhz MicroSPARC II controller - 6 SCSI-2 buses - 42x4 GB Disks (6048) GB total) - 2x25 Mbps FibreChannel (FC) Interfaces - Configured for RAID 10

Figure 5-12. Centralized Process/Centralized Data Architecture Components (cont'd)

[15] According to Netscape's *SuiteSpot 3.0 Evaluation Guide*, provided adequate resources, the Netscape's Directory Server is capable of supporting up to 4 million directory entries and as many as 300,000 queries per hour.

[16] According to DAZEL Corporation, one full time staff equivalent is typically required to manage and administer 50 printers; however, using the DAZEL Output Server, Barnet Bank reports that 5 help desk staff are capable of servicing and administering more than 800 printers.

[17] According to Sun Microsystems, Inc., SPARCstorage Array technology is widely used. For example, Shomega Limited uses SPARCstorage Array 100s and 200s to provide a data warehousing and imaging solution that stores more than 250,000 full color images on more than 730 GB of disk space.

Reference	Vendor/Product	Description	Quantity	Hardware Specification
	SPARCstorage Array 214 RSM	Data Center Class Storage Array Subsystem (Applications)	3	- 1x110 Mhz MicroSPARC II controller - 6 SCSI-2 buses - 18x4 GB Disks (210 GB total) - 2x25 Mbps FibreChannel (FC) Interfaces - Configured for RAID 5
3	SPARCstorage Array 219	Electronic Mail Storage Array Subsystem	3	- 1x110 Mhz MicroSPARC II controller - 6 SCSI-2 buses - 18x9 GB Disks (478 GB total) - 2x25 Mbps FibreChannel (FC) Interfaces - Configured for RAID 5

Figure 5-12. Centralized Process/Centralized Data Architecture Components (cont'd)

Figure 5.13, details storage capacity calculations for the centralized process/centralized data candidate architecture. These calculations were used to estimate required storage capacity.

Service	Raw Data Calculation[A]	Total Capacity	RAID Configuration	RAID Calculation[B]	Instances	Total RAID Capacity	Disk Configuration[C]
Telephony Service, Web Service, EDI Service, Output Management Service, OLTP Monitor Service, Application Software Services	168GB	168 GB	RAID 5	25% increase	1	210GB	3x18x4
Electronic Mail	378GB	378GB	RAID 5	25% increase	1	478GB	3x18x9
OLTP Database Service	1490GB	2980GB	RAID 10	100% increase	1	5960GB	36x42x4
Total Required Storage Capacity Requirements						6648GB	

Figure 5-13. Centralized Process/Centralized Data Architecture Storage Capacity

[A] According to the reference text *Capacity and Configuration Planning for Solaris Servers*, unless more accurate information is available, DBMS' should be configured with approximately twice as much disk space as raw data to account for additional database structures, such as table indexes, temporary tables, log files, archives, etc. Taking this into account, the "Raw Data Calculation" represents the disk space available for raw data, assuming the remaining available disk space will be consumed by other database structures.

[B] Where RAID 1+0 is used, required disk capacity is doubled to accommodate disk mirroring (shadowing). Where RAID 5 is used it is assumed that roughly 20 percent of disk capacity will be consumed by parity data, which is used to recover from failures. As a result, required disk space is increased by 25 percent.

[C] RAID configurations are expressed as AxDxC, where A equals number of storage array subsystems, D equals number disks within each subsystem, and C equals the capacity, in gigabytes, of each disks within the subsystem. Where "0" storage array subsystems are denoted, disk storage is provided by a single disk. This configuration provides the "Total RAID Capacity" required by the services.

5.4.2 Distributed Process/Centralized Data Architecture

This architecture includes centralized data management resources, which facilitate the execution of all data management processing (transaction processing and decision support) from a single computing resource. Unlike the fully centralized architecture, this architecture allows for the distribution of application and presentation processing resources and activities. This architecture is typically characterized as:

- **Performance constrained** – Common processing resource required to simultaneously support data management operations associated with decision support and transaction processing activities.
- **Tightly coupled** – Autonomous and disparate operations constrained by reliance on centralized data resource for decision support and transaction processing activities.
- **Vulnerable** – All data dependent system operations can be catastrophically affected by single resource failure.
- **Specialized** – Provides opportunity to use specialized technology resources, which are specifically engineered for supporting and providing particular services.
- **Flexible** –System growth and service support not constrained by single hardware environment and limited vendor or third party product offerings.

With these characteristics in mind, the architecture for Project EASI/ED was designed to mitigate weaknesses and leverage strengths intrinsic to distributed process/centralized data architectures. This architecture is detailed in Figure 5.14, shown on the following page.

Figure 5.14 describes interfaces between hardware components and complements the software interface descriptions provided in Appendix C.

The distributed process/centralized data candidate framework architecture is organized in terms of four technical "classifications." These classifications are:

- **Electronic Commerce Resources** – provide the software and hardware technologies required to support telephony, Web, and EDI-based processing. Specifically, these technologies allow users to access and invoke system services, such as those provided by transaction processing and decision support technologies, through commonly available Web browsers and telephone technology. Additionally, these technologies allow trading partners to electronically communicate via X12 compliant transaction sets.

- **Infrastructure Resources –** provide the software and hardware technologies required to support the "Centralized Software Management and Dynamic Deployment" strategy described in Subsection 5.3.7. These technologies support DNS name services and provide the NFS services required to support the AutoClient server configurations used for the Web, transaction processing, and application display servers. These services largely transparent to the user, but simplify system administration, maintenance, and operations.

- **Communication Resources** – provide the software and hardware technologies required to support E-mail, printer, and facsimile services. Additionally, these technologies provide the "client-side" remote application display services used to invoke and execute system services via LAN/WAN-based X Window System clients.

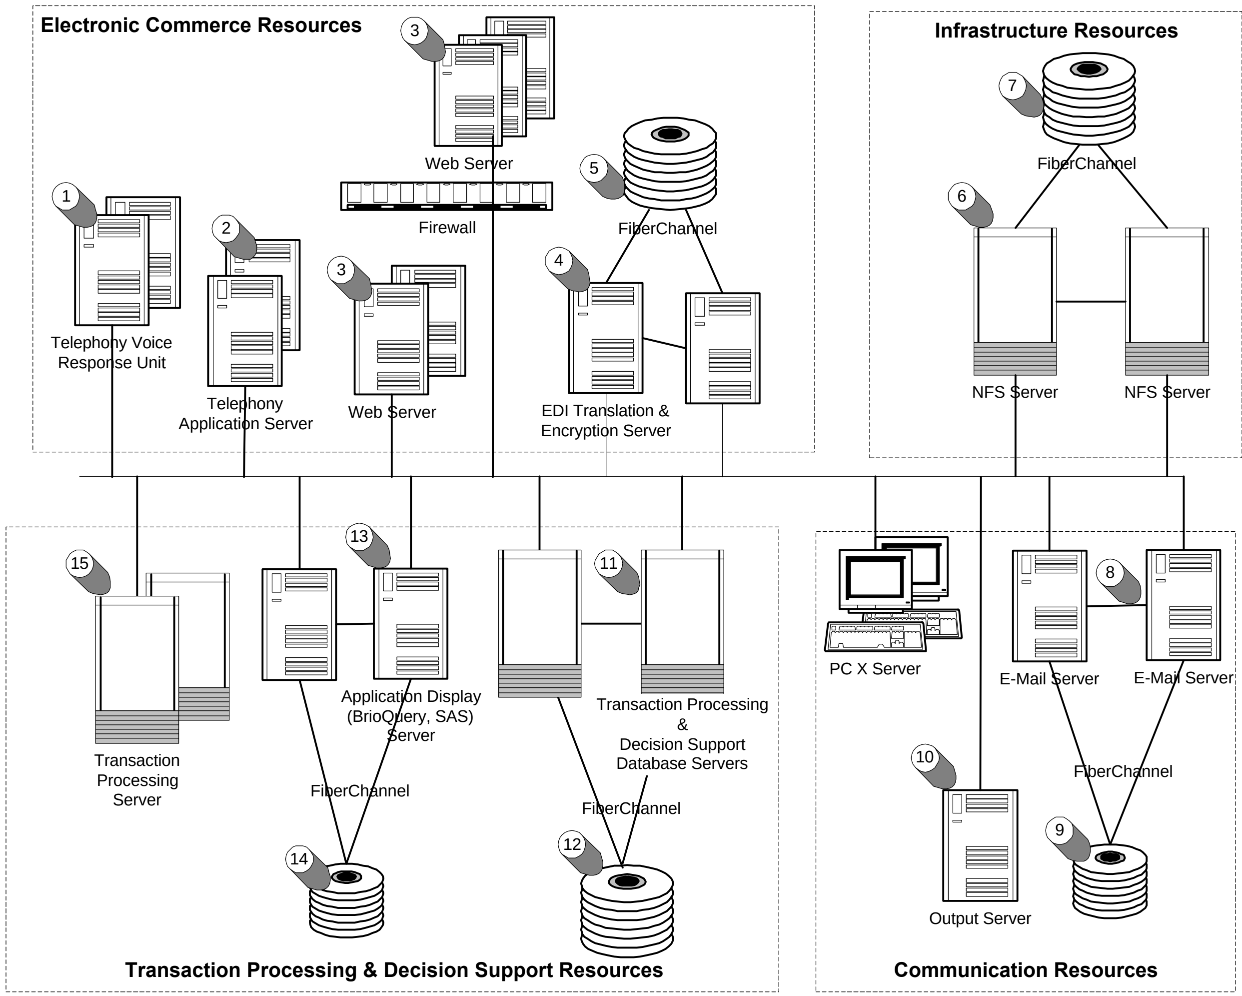

Figure 5-14. Distributed Process/Centralized Data Architecture

- **Transaction Processing and Decision Support** – provide the software and hardware technologies required to support on-line and batch-oriented transaction processing and decision support operations. Transaction processing servers are used to process “component-based” transaction processing services. These services are invoked in batch, via electronic commerce technologies (Telephony and Web-based) and by communication resources (X Window System clients).

Because of architectural constraints (i.e. centralized data), data warehousing services are not available within this architecture. As a result, the database servers host a common relational database, used for both transaction processing and decision support operations.

The application display servers are required to support the “Centralized Execution and Remote Display of Application Clients” – a strategy described in Subsection 5.3.8. The application display servers support execution of multiple client application software instantiations, which are remotely displayed on remote X Window System communication resources.

Figure 5.15, describes the hardware and software components that comprise the distributed process/centralized data candidate framework architecture.

Reference	Vendor/Product	Description	Quantity	Hardware Specification
1	VoiceTek, Inc VTK 3000	Interactive Voice Response Unit Server	3	- 1 Pentium processor - 166Mhz clock speed - 32MB RAM - 4 GB of Disk - SCO UNIX 5.0.2 - T1 trunk interface - 100BaseT LAN interface
1	VoiceTek, Inc. Generations Telephony Server Platform v4.0	Telephony – Interactive Voice Response Unit Software	3	
2	Sun Microsystems, Inc. Ultra Enterprise 1	Enterprise Class Telephony Application Hardware Server	6	- 1 UltraSPARC processor - 167Mhz clock speed - 64 MB RAM - 2 GB Disk - Solaris 2.6 - 100BaseT LAN Interface
	VoiceTek, Inc. Generations Runtime Server Platform v4.0	Telephony – Runtime Application Server	6	
3	Sun Microsystems, Inc. Ultra Enterprise 1[18]	Web Hardware Server	3	- 1 UltraSPARC processor - 167Mhz clock speed - 64 MB RAM - 2 GB Disk - Solaris 2.6 - 100BaseT LAN Interface

Figure 5.15. Distributed Process/Centralized Data Architecture Components

[18] *Configuration and Capacity Planning for Solaris Servers (a reference text)* reports that multiprocessor scalability of Solaris Web performance is substantially lower than for most other operations – scaling well to only two processors. As a result, busy web sites must resort to server farms and techniques such as round-robin DNS, rather than large multiprocessors. These strategies are followed throughout all of the distributed architectures.

Reference	Vendor/Product	Description	Quantity	Hardware Specification
4	Sun Microsystems, Inc Ultra Enterprise 2	Highly Available, Clustered Enterprise Class EDI Hardware Server	2	- 1 UltraSPARC processor - 167 Mhz clock speed - 256 MB RAM - Solaris 2.6 - 100BaseT LAN Interface
	Solstice HA v1.3	High Availability Hardware Monitoring and Fault Detection Software	1	
	Solstice: DiskSuite v4.0	High Availability Volume Management Software	1	
	Premenos Corporation Templar Enterprise	EDI Security and SMTP/MIME Delivery Software	1	
	Premenos Corporation EDI/Open	EDI X12 Translation Software	1	
	Oracle Corporation Oracle 8 Server	Relational Database Management System Software	1	
5	SPARCstorage Array 112 RSM	Storage Array Subsystem	1	- 1x110 Mhz MicroSPARC II controller - 6 SCSI-2 buses - 13x2 GB Disks (26 GB total) - 2x25 Mbps FibreChannel (FC) Interfaces - Configured for RAID 5
6	Sun Microsystems, Inc. Ultra Enterprise 2[19]	Highly Available, Clustered Enterprise Class NFS Hardware Servers	2	- 1 UltraSPARC processor - 167 Mhz clock speed - 128 MB RAM - Solaris 2.6 - 100BaseT LAN Interface
	Netscape Communication Corporation Enterprise Server v3.0	Web Server Software	3	
	Netscape Communication Corporation Directory Server	Web Directory Server Software	1	
	SAS Institute, Inc SAS	Statistical and Quantitative Analysis Software	1 (site license)	
	Solstice HA v1.3	High Availability Hardware Monitoring and Fault Detection Software	1	
	Solstice: DiskSuite v4.0	High Availability Volume Management Software	1	

Figure 5.15. Distributed Process/Centralized Data Architecture Components (cont'd)

[19] According to *Configuration and Capacity Planning for Solaris Server* (a reference text) the Ultra 1/170 can handle 65-70 CacheOS NFS clients even in the worst case of all booting simultaneously. Although the architecture does not include 65-70 CacheOS, the Ultra Enterprise 2 (which is significantly more powerful that the Ultra 1/170) is being used as the NFS server, as an Ultra Enterprise class server is required in order to implement high availability clustered configurations. This strategy is followed throughout all of the architectures.

Reference	Vendor/Product	Description	Quantity	Hardware Specification
6	HA NFS	Highly Available Network File System Software	1	
	RRDNS	Primary Round-Robin Domain Name Server Software	1	
7	SPARCstorage Array 112 RSM	Storage Array Subsystem	2	- 1x110 Mhz MicroSPARC II controller - 6 SCSI-2 buses - 19x2 GB Disks (76 GB total) - 2x25 Mbps FibreChannel (FC) Interfaces - Configured for RAID 5
8	Sun Microsystems, Inc. Ultra Enterprise 2[20]	Highly Available, Clustered Enterprise Class Electronic Mail Hardware Servers	2	- 1 UltraSPARC processor - 167 Mhz clock speed - 450 MB RAM - Solaris 2.6 - 100BaseT LAN Interface
	Sun Microsystems, Inc. Solstice Internet Mail Server v2.0	Internet Electronic Mail Server Software	1	
	Solstice HA v1.3	High Availability Hardware Monitoring and Fault Detection Software	1	
	Solstice: DiskSuite v4.0	High Availability Volume Management Software	1	
	RRDNS	Secondary Round-Robin Domain Name Server Software	1	
9	SPARCstorage Array 219 RSM	Storage Array Subsystem	3	- 1x110 Mhz MicroSPARC II controller - 6 SCSI-2 buses - 18x9 GB Disks (478 GB total) - 2x25 Mbps FibreChannel (FC) Interfaces - Configured for RAID 10
10	Sun Microsystems, Inc. Ultra Enterprise 1	Output Management Hardware Server	1	- 1 UltraSPARC processor - 167Mhz clock speed - 64 MB RAM - 4 GB Disk - Solaris 2.6 - 100BaseT LAN Interface
	DAZEL Corporation DAZEL Output Server	Output Management System Software	1	

Figure 5.15. Distributed Process/Centralized Data Architecture Components (cont'd)

[20] Sun Microsystems, Inc. reports that the Solstice Internet Mail Server is capable of simultaneously supporting more than 1,800 active IMAP users on a single Ultra Enterprise 1 server. Sun is currently using this architecture providing access to 20,000 users processing 4-5 million messages per day. The more powerful Ultra Enterprise 2 server is required within the candidate architectures in order to implement a highly available, clustered electronic mail solution.

Reference	Vendor/Product	Description	Quantity	Hardware Specification
11	Sun Microsystems, Inc. Ultra Enterprise 6000[21] [22]	Highly Available, Clustered Enterprise Class Database Management Hardware Servers	1	- 24 UltraSPARC processors - 250 Mhz clock speed - 6 GB RAM - Solaris 2.6 - 100BaseT LAN Interface
	Oracle Corporation Oracle 8 Parallel Server	Parallel Processing Relational Database Management System Software	1	
	Solstice PDB v1.2	High Availability, Parallel Processing Hardware Monitoring and Fault Detection Software	1	
	BEA Systems, Inc. HA Extensions for BEA Tuxedo	High Availability, Hardware Monitoring and Fault Detection Software for Tuxedo	1	
	Solstice: DiskSuite v4.0	High Availability Volume Management Software	1	
12	SPARCstorage Array 214 RSM[23]	Storage Array Subsystem	36	- 1x110 Mhz MicroSPARC II controller - 6 SCSI-2 buses - 42x4 GB Disks (7392 GB total) - 2x25 Mbps FibreChannel (FC) Interfaces - Configured for RAID 10
13	Sun Microsystems, Inc. Ultra Enterprise 2	High Availability Cluster Enterprise Class Application Display Hardware Server	2	- 2 UltraSPARC processors - 300 Mhz clock speed - 1 GB RAM - Solaris 2.6 - 100BaseT LAN Interface
14	SPARCstorage Array 112 RSM	Storage Array Subsystem	2	- 1x110 Mhz MicroSPARC II controller - 6 SCSI-2 buses - 19x2 GB Disks (76 GB total) - 2x25 Mbps FibreChannel (FC) Interfaces - Configured for RAID 10

Figure 5.15. Distributed Process/Centralized Data Architecture Components (cont'd)

[21] Sun Microsystems, Inc., reports that, Pfizer, Inc. is successfully using an Ultra Enterprise 6000 cluster to provide manufacturing plants with continuous 24x7 access to Oracle database resources.
[22] According to TPC-C benchmark results, a configuration including the BEA OLTP Monitor and Oracle 8 running on an Ultra Enterprise 6000 is capable of simultaneously supporting 26,000 users and delivering a throughput equivalent to 31,147.04 transactions per minute.
[23] Equivalent storage capacity can be provided with the eight 219 RSM SPARCstorage Arrays, which individually provide over 370 GB of disk space. However, the smaller 214 model will allow I/O to be distributed across 18 controllers, which will compliment the disk stripping services provided by the RAID configuration and improve concurrent (parallel) disk I/O. This strategy is followed throughout all of the architectures.

Reference	Vendor/Product	Description	Quantity	Hardware Specification
15	Sun Microsystems, Inc Ultra Enterprise 1	Batch and Online Transaction Processing Monitor Hardware Server	15	- 1 UltraSPARC processor - 167 Mhz clock speed - 512 MB RAM - 2.1 GB Disk - Solaris 2.6 - 100BaseT LAN Interface
	BEA Systems, Inc. BEA Tuxedo v6.x	Online Transaction Processing Monitor	15	
	BEA Systems, Inc. HA Extensions – HA Client Wrapper	High Availability, In-flight transaction recovery and resubmission software	15	
	BEA Systems, Inc. BEA Jolt v1.1	Web Gateway for Online Transaction Processing Monitor Services	15	

Figure 5.15. Distributed Process/Centralized Data Architecture Components (cont'd)

Figure 5.16, details storage capacity calculations for the distributed process/centralized data candidate architecture. These calculations were used to estimate required storage capacity.

Service	Raw Data Calculation[A]	Total Capacity	RAID Configuration	RAID Calculation[B]	Instances	Total RAID Capacity	Disk Configuration[C]
Telephony Service	2GB	2 GB	NA	NA	6	12GB	0x1x2
Web Service	2GB	2GB	NA	NA	3	6GB	0x1x2
EDI Service	20GB	20GB	RAID 5	25% increase	1	25GB	1x13x2
NFS Service	60GB	60GB	RAID 5	25% increase	1	76GB	2x19x2
Application Service	60GB	60GB	RAID 5	25% increase	1	76GB	2x19x2
Electronic Mail Service	378GB	378GB	RAID 5	25% increase	1	478GB	3x18x9
Output Management Service	4GB	4GB	NA	NA	1	4GB	0x1x2
OLTP Database Service	1490GB	2980GB	RAID 10	100% increase	1	5960GB	36x42x4
OLTP Monitor Service	2GB	2GB	NA	NA	15	30GB	0x1x2
Total Required Storage Capacity Requirements						6667GB	

Figure 5.16. Distributed Process/Centralized Data Architecture Storage Capacity

[A]According to the reference text *Capacity and Configuration Planning for Solaris Servers*, unless more accurate information is available, DBMS' should be configured with approximately twice as much disk space as raw data to account for additional database structures, such as table indexes, temporary tables, log files, archives, etc. Taking this into account, the "Raw Data Calculation" represents the disk space available for raw data, assuming the remaining available disk space will be consumed by other database structures.

[B] Where RAID 1+0 is used, required disk capacity is doubled to accommodate disk mirroring (shadowing). Where RAID 5 is used it is assumed that roughly 20 percent of disk capacity will be consumed by parity data, which is used to recover from failures. As a result, required disk space is increased by 25 percent.

[C] RAID configurations are expressed as AxDxC, where A equals number of storage array subsystems, D equals number disks within each subsystem, and C equals the capacity, in gigabytes, of each disks within the subsystem. Where "0" storage array subsystems are denoted, disk storage is provided by a single disk.

5.4.3 Distributed Process/Replicated Data for Consolidation Architecture

This architecture includes centralized decision support data management resources. However, unlike architectures that employ a fully centralized data management strategy, this architecture allows for the distribution of resources associated with transaction processing data management, and application and presentation services.

The replication for data consolidation configuration facilitates collection of data from multiple primary sites, each of which support transaction processing activities. The data consolidation configuration is often useful in situations where data may need to be regularly aggregated and reviewed, but distributed components need to be able to work independently or off-line.

This architecture is typically characterized as:

- **Scaleable** – Common processing resources are not required to simultaneously support data management operations associated with decision support and transaction processing activities. As a result, specialized and heterogeneous resources can be employed to deliver required services.
- **Loosely coupled** – Autonomous operations do not rely on a single, centralized data resource for decision support and transaction processing activities. Rather, disparate activities, such as those associated with decision support and transaction processing, leverage specialized and largely independent resources.
- **Complex** – Disparate technologies mandate specialized skills and increase complexity of management, operation, and administration of heterogeneous technologies.
- **Fault tolerant** – Distributed data management and application processing allow for increased fault tolerance, as all system services and associated processing do not rely on any single resource.
- **Specialized –** Provides opportunity to use specialized technology resources, which are specifically engineered for supporting and providing particular services.
- **Flexible** –System growth, and service support not constrained by single hardware environment and limited vendor or third party product offerings.

With these characteristics in mind, the architecture for Project EASI/ED was designed to mitigate weaknesses and leverage strengths intrinsic to centralized process/replication for data consolidation. This architecture is detailed in Figure 5.17, shown on the following page.

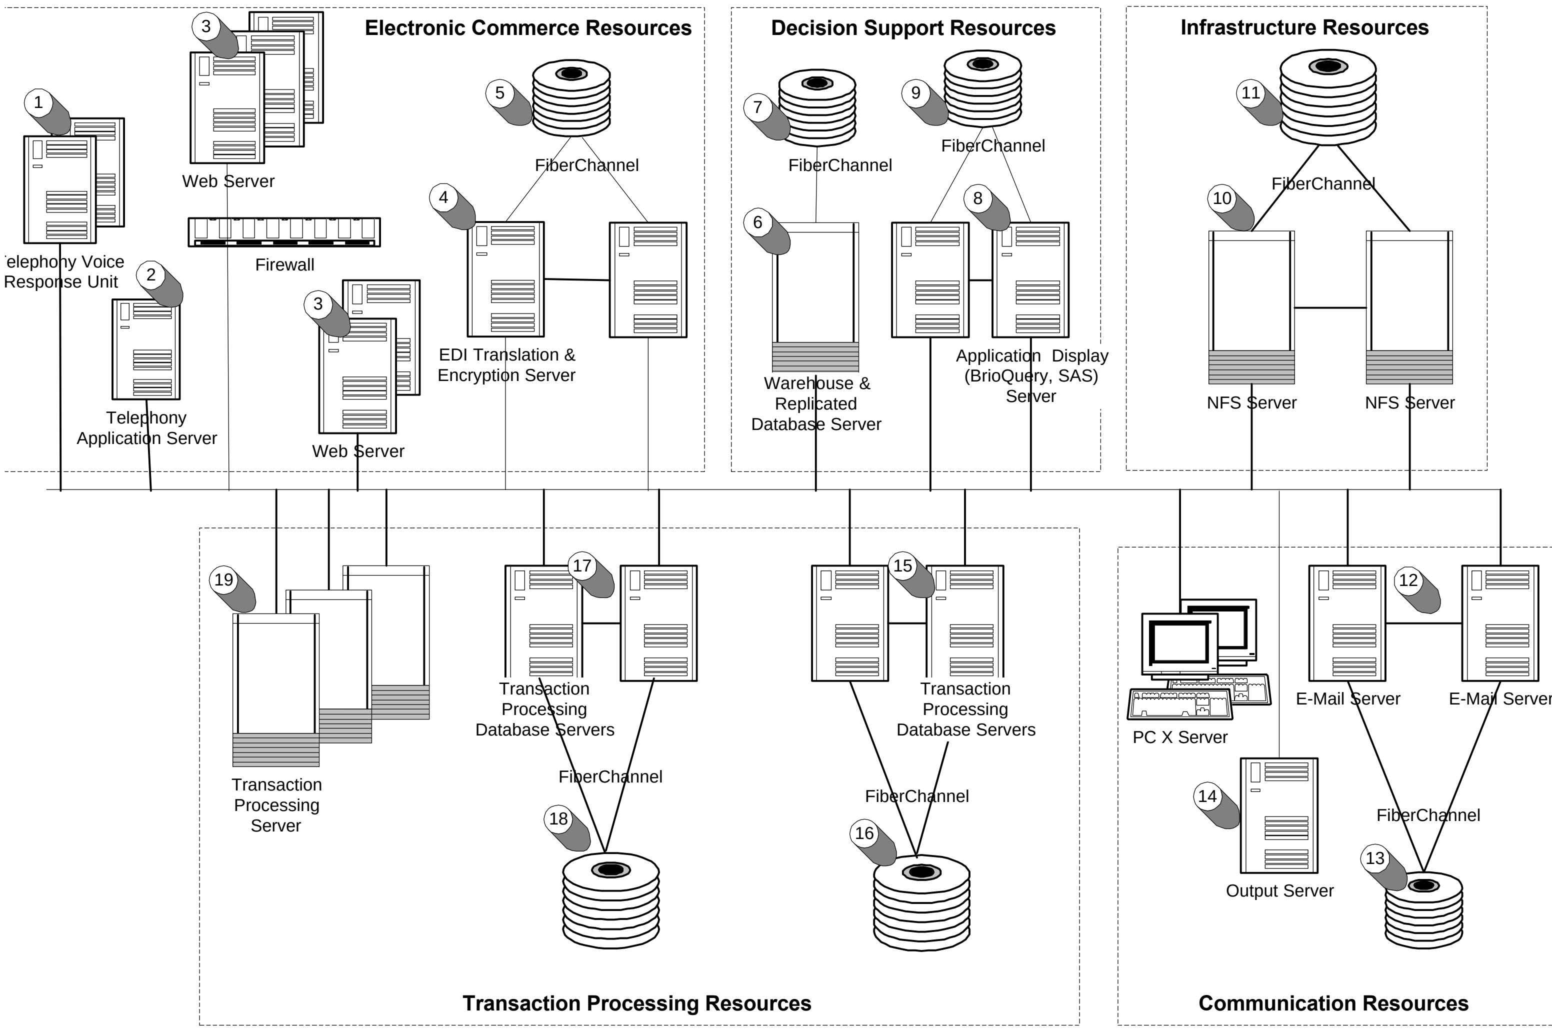

Figure 5.17. Distributed Process/ Replication for Data Consolidation Architecture

Figure 5.17 describes interfaces between hardware components and complements the software interface descriptions provided in Appendix C.

The distributed process/replicated data for consolidation architecture is organized in terms of five technical "classifications." These classifications are:

- **Electronic Commerce Resources** – provide the software and hardware technologies required to support telephony, Web, and EDI-based processing. Specifically, these technologies allow users to access and invoke system services, such as those provided by transaction processing and decision support technologies, through commonly available Web browsers and telephone technology. Additionally, these technologies allow trading partners to electronically communicate via X12 compliant transaction sets.

- **Infrastructure Resources –** provide the software and hardware technologies required to support the "Centralized Software Management and Dynamic Deployment" strategy described in Subsection 5.3.7. These technologies support DNS name services and provide the NFS services required to support the AutoClient server configurations that are used for Web, transaction processing, and application display servers. These services are largely transparent to users, but simplify system administration, maintenance, and operations.

- **Communication Resources –** provide the software and hardware technologies required to support E-mail, printer, and facsimile services. Additionally, these technologies provide the "client-side" remote application display services that are used to invoke and execute system services via LAN/WAN-based X Window System clients.

- **Transaction Processing Resources** – provide the software and hardware technologies required to support on-line and batch-oriented transaction processing operations. Transaction processing servers are used to provide "component-based" transaction processing services. These services are invoked in batch, via electronic commerce technologies (telephony and Web-based) and by communication resources (X Window System clients).

  Unlike previously described architectures, multiple database servers are exclusively dedicated to supporting transaction-processing operations. Decision support operations are executed against other resources and are not supported by the transaction processing database servers.

- **Decision Support Resources –** provide the software and hardware technologies required by decision support operations. The data warehousing and replicated database server (a single hardware platform) serves a "read-only" replica subset of the transaction processing database, as well as a data warehouse. The warehouse is loaded from the transaction processing databases and includes derived and aggregated data.

  Application display servers are required to support the "Centralized Execution and Remote Display of Application Clients" – a strategy described in Subsection 5.3.8. Application display servers support execution of multiple client application software instantiations, which are remotely displayed on remote X Window System communication resources. These clients primarily include decision support tools such as BrioQuery and SAS. However, X Window System transaction process application clients also leverage this server's resources.

The following table, Figure 5.18, describes the hardware and software components that comprise the distributed process/replicated data for consolidation candidate framework architecture.

Reference	Vendor/Product	Description	Quantity	Hardware Specification
1	VoiceTek, Inc VTK 3000	Interactive Voice Response Unit Server	3	- 1 Pentium processor - 166Mhz clock speed - 32MB RAM - 4 GB of Disk - SCO UNIX 5.0.2 - T1 trunk interface - 100BaseT LAN interface
	VoiceTek, Inc. Generations Telephony Server Platform v4.0	Telephony – Interactive Voice Response Unit Software	3	
2	Sun Microsystems, Inc. Ultra Enterprise 1[24]	Enterprise Class Telephony Application Hardware Server	6	- 1 UltraSPARC processor - 167Mhz clock speed - 64 MB RAM - 2 GB Disk - Solaris 2.6 - 100BaseT LAN Interface
	VoiceTek, Inc. Generations Runtime Server Platform v4.0	Telephony – Runtime Application Server	6	
3	Sun Microsystems, Inc. Ultra Enterprise 1[25]	Web Hardware Server	3	- 1 UltraSPARC processor - 167Mhz clock speed - 64 MB RAM - 2 GB Disk - Solaris 2.6 - 100BaseT LAN Interface
4	Sun Microsystems, Inc Ultra Enterprise 2	Highly Available, Clustered Enterprise Class EDI Hardware Server	2	- 1 UltraSPARC processor - 167 Mhz clock speed - 256 MB RAM - Solaris 2.6 - 100BaseT LAN Interface
	Solstice HA v1.3	High Availability Hardware Monitoring and Fault Detection Software	1	
	Solstice: DiskSuite v4.0	High Availability Volume Management Software	1	
	Premenos Corporation Templar Enterprise	EDI Security and SMTP/MIME Delivery Software	1	
	Premenos Corporation EDI/Open	EDI X12 Translation Software	1	
	Oracle Corporation Oracle 8 Server	Relational Database Management System Software	1	

Figure 5.18. Distributed Process/Replicated Data for Consolidation Architecture Components

[24] Sun Microsystems Inc. reports that Scottish Power, Plc., provides call center telephony support for more that 500 concurrent users on two SPARCserver 1000E's via telephony.

[25] According to *Configuration and Capacity Planning for Solaris Server* (a reference text) a SPARCserver 1000E (as configured) is capable of supporting 7.7 million hits per day (assuming no other load is placed on the server). This allows for up to 90 sustained connections per second. Working together the three Web servers recommended within the distributed architectures will be capable of supporting a 23.1 million hits a day. These servers will be access by Internet users via a T3 facility, which is capable of supporting 23.2 million hits per day (assuming the hits are 5Kb in size on average). This strategy is followed throughout all of the architectures.

Reference	Vendor/Product	Description	Quantity	Hardware Specification
5	SPARCstorage Array 112 RSM	Storage Array Subsystem	1	- 1x110 Mhz MicroSPARC II controller - 6 SCSI-2 buses - 13x2 GB Disks (26 GB total) - 2x25 Mbps FibreChannel (FC) Interfaces - Configured for RAID 5
6	Sun Microsystems, Inc. Ultra Enterprise 4000	Enterprise Class Decision Support Data Management Server	1	- 12 UltraSPARC processors - 250 Mhz clock speed - 3 GB RAM - Solaris 2.6 - 100BaseT LAN Interface
	Solstice: DiskSuite v4.0	High Availability Volume Management Software	1	
	Red Brick Systems, Inc. Red Brick Warehouse v5.0[26]	Data Warehouse Relation al Database Management system Software	1	
	Oracle Corporation Oracle 8 Server	Relational Database Management System Software	1	
7	SPARCstorage Array 214 RSM	Storage Array Subsystem	28	- 1x110 Mhz MicroSPARC II controller - 6 SCSI-2 buses - 42x4 GB Disks (4620GB total) - 2x25 Mbps FibreChannel (FC) Interfaces - Configured for RAID 5
8	Sun Microsystems, Inc. Ultra Enterprise 2	High Availability Cluster Enterprise Class Application Display Hardware Server	2	- 2 UltraSPARC processors - 300 Mhz clock speed - 1 GB RAM - Solaris 2.6 - 100BaseT LAN Interface
9	SPARCstorage Array 112 RSM	Storage Array Subsystem	2	- 1x110 Mhz MicroSPARC II controller - 6 SCSI-2 buses - 19x2 GB Disks (76 GB total) - 2x25 Mbps FibreChannel (FC) Interfaces - Configured for RAID 10
10	Sun Microsystems, Inc. Ultra Enterprise 2	Highly Available, Clustered Enterprise Class NFS Hardware Servers	2	- 1 UltraSPARC processor - 167 Mhz clock speed - 128 MB RAM - Solaris 2.6 - 100BaseT LAN Interface
	Netscape Communication Corporation Enterprise Server v3.0[27]	Web Server Software	3	

Figure 5.18. Distributed Process/Replicated Data for Consolidation Architecture Components (cont'd)

[26] Red Brick Systems, Inc. reports that a leading retail customer is loading 35,000,000 rows into a 300 GB data warehouse every night. Using 6 CPUs, the complete load process is taking an average of 35 minutes each night. Red Brick's attributes this performance to the product's parallel load feature, which is capable of sustaining a rate of over 150,000 records per minute per CPU.

[27] Return On Investment (ROI) case studies, conducted by International Data Corporation, indicate that Netscape Intranet users, such as Amdahl, Booz-Allen and Hamilton, and Silicon Graphics, find the typical ROI to be over 1,000 percent.

Reference	Vendor/Product	Description	Quantity	Hardware Specification
10	Brio Technology, Inc. BrioQuery.Insight	Online Analytical Processing Plug-in for the Netscape Enterprise Server	3	
	Brio Technology, Inc. BrioQuery.Quickview	Online Analytical Processing Plug-in for the Netscape Enterprise Server	3	
	Netscape Communication Corporation Directory Server	Web Directory Server Software	1	
	Brio Technology, Inc. BrioServer	Data Warehouse Extract Utility for Online Analytical Processing	1 (site license)	
	SAS Institute, Inc SAS	Statistical and Quantitative Analysis Software	1 (site license)	
	Solstice HA v1.3	High Availability Hardware Monitoring and Fault Detection Software	1	
	Solstice: DiskSuite v4.0	High Availability Volume Management Software	1	
	HA NFS	Highly Available Network File System Software	1	
	RRDNS	Primary Round-Robin Domain Name Server Software	1	
11	SPARCstorage Array 112 RSM	Storage Array Subsystem	2	- 1x110 Mhz MicroSPARC II controller - 6 SCSI-2 buses - 19x2 GB Disks (76 GB total) - 2x25 Mbps FibreChannel (FC) Interfaces - Configured for RAID 5
12	Sun Microsystems, Inc. Ultra Enterprise 2	Highly Available, Clustered Enterprise Class Electronic Mail Hardware Servers	2	- 1 UltraSPARC processor - 167 Mhz clock speed - 450 MB RAM - Solaris 2.6 - 100BaseT LAN Interface
	Sun Microsystems, Inc. Solstice Internet Mail Server v2.0	Internet Electronic Mail Server Software	1	
	Solstice HA v1.3	High Availability Hardware Monitoring and Fault Detection Software	1	
	Solstice: DiskSuite v4.0	High Availability Volume Management Software	1	
	RRDNS	Secondary Round-Robin Domain Name Server Software	1	

Figure 5.18. Distributed Process/Replicated Data for Consolidation Architecture Components (cont'd)

Reference	Vendor/Product	Description	Quantity	Hardware Specification
12	SPARCstorage Array 219 RSM	Storage Array Subsystem	3	- 1x110 Mhz MicroSPARC II controller - 6 SCSI-2 buses - 18x9 GB Disks (486 GB total) - 2x25 Mbps FibreChannel (FC) Interfaces - Configured for RAID 10
13	Sun Microsystems, Inc. Ultra Enterprise 1	Output Management Hardware Server	1	- 1 UltraSPARC processor - 167Mhz clock speed - 64 MB RAM - 4 GB Disk - Solaris 2.6 - 100BaseT LAN Interface
14	DAZEL Corporation DAZEL Output Server	Output Management System Software	1	
15	Sun Microsystems, Inc. Ultra Enterprise 4000	Highly Available, Clustered Enterprise Class Database Management Hardware Servers	2	- 12 UltraSPARC processors - 250 Mhz clock speed - 3 GB RAM - Solaris 2.6 - 100BaseT LAN Interface
	Oracle Corporation Oracle 8 Parallel Server	Parallel Processing Relational Database Management System Software	1	
	Solstice PDB v1.2	High Availability, Parallel Processing Hardware Monitoring and Fault Detection Software	1	
	Solstice: DiskSuite v4.0	High Availability Volume Management Software	1	
16	SPARCstorage Array 214 RSM[28]	Storage Array Subsystem	22	- 1x110 Mhz MicroSPARC II controller - 6 SCSI-2 buses - 42x4 GB Disks (3696 GB total) - 2x25 Mbps FibreChannel (FC) Interfaces - Configured for RAID 10

Figure 5.18. Distributed Process/Replicated Data for Consolidation Architecture Components (cont'd)

[28] Oracle allows DBMS files to be stored either in raw disks or in standard UNIX file systems. However, storage in the file system is typically less efficient. According to *Configuration and Capacity Planning for Solaris Servers* (a reference text) file system DBMS storage is usually 5 percent, but never more than 15 percent, less efficient – since an additional layer of system software must be traversed for every DBMS disk access. Additionally, file systems consume roughly 10 percent of available disk space with overhead data about the filesystem. For these reasons, DBMS files will be stored on raw disk partitions and not within filesystems. This strategy is followed throughout all of the architectures.

Reference	Vendor/Product	Description	Quantity	Hardware Specification
17	Sun Microsystems, Inc. Ultra Enterprise 4000	Highly Available, Clustered Enterprise Class Database Management Hardware Servers	2	- 12 UltraSPARC processors - 250 Mhz clock speed - 3 GB RAM - Solaris 2.6 - 100BaseT LAN Interface
	Oracle Corporation Oracle 8 Parallel Server	Parallel Processing Relational Database Management System Software	1	
	Solstice PDB v1.2	High Availability, Parallel Processing Hardware Monitoring and Fault Detection Software	1	
	Solstice: DiskSuite v4.0	High Availability Volume Management Software	1	
18	SPARCstorage Array 214 RSM[29]	Storage Array Subsystem	22	- 1x110 Mhz MicroSPARC II controller - 6 SCSI-2 buses - 42x4 GB Disks (3696 GB total) - 2x25 Mbps FibreChannel (FC) Interfaces - Configured for RAID 10
19	Sun Microsystems, Inc Ultra Enterprise 1	Batch and Online Transaction Processing Monitor Software	15	- 1 UltraSPARC processor - 167 Mhz clock speed - 512 MB RAM - 2.1 GB Disk - Solaris 2.6 - 100BaseT LAN Interface
	BEA Systems, Inc. BEA Tuxedo v6.x	Online Transaction Processing Monitor	15	
	BEA Systems, Inc. HA Extensions – HA Client Wrapper	High Availability, In-flight transaction recovery and resubmission software	15	
	BEA Systems, Inc. BEA Jolt v1.1	Web Gateway for Online Transaction Processing Monitor Services	15	

Figure 5.18. Distributed Process/Replicated Data for Consolidation Architecture Components (cont'd)

[29] DBMS systems typically use a variety of data/storage structures including data tables, indices, temporary tables, transaction logs, rollback buffers, and system configuration files. To minimize I/O contention, the operating system and swap files, data, indexes, log files, and rollback buffers should be distributed across different disks or disk sets. This strategy is followed throughout all of the architectures.

Figure 5.19, details storage capacity calculations for the distributed process/replication for data consolidation candidate framework architecture. These calculations were used to estimate required storage capacity.

Service	Raw Data Calculation[A]	Total Capacity	RAID Configuration	RAID Calculation[B]	Instances	Total RAID Capacity	Disk Configuration[C]
Telephony Service	2GB	2 GB	NA	NA	6	12GB	0x1x2
Web Service	2GB	2GB	NA	NA	3	6GB	0x1x2
EDI Service	20GB	20GB	RAID 5	25% increase	1	25GB	1x13x2
Decision Support RDBMS and Data Warehouse Service	1490GB	2980GB	RAID 5	25% increase	1	3725GB	23x42x4
NFS Service	60GB	60GB	RAID 5	25% increase	1	76GB	2x19x2
Application Services	60GB	60GB	RAID 5	25% increase	1	76GB	2x19x2
Electronic Mail Service	378GB	378GB	RAID 5	25% increase	1	478GB	3x18x9
Output Management Service	4GB	4GB	NA	NA	1	4GB	0x1x2
OLTP Database Service (Cluster 1)	570GB	1140GB	RAID 10	100% increase	1	2280GB	14x42x4
OLTP Database Service (Cluster 2)	570GB	1140GB	RAID 10	100% increase	1	2280GB	14x42x4
OLTP Monitor Service	2GB	2GB	NA	NA	15	30GB	0x1x2
Total Required Storage Capacity Requirements						8992GB	

Figure 5.19. Distributed Process/Replicated Data for Consolidation Architecture Storage Capacity

[A] According to the reference text *Capacity and Configuration Planning for Solaris Servers*, unless more accurate information is available, DBMS' should be configured with approximately twice as much disk space as raw data to account for additional database structures, such as table indexes, temporary tables, log files, archives, etc. Taking this into account, the "Raw Data Calculation" represents the disk space available for raw data, assuming the remaining available disk space will be consumed by other database structures.

[B] Where RAID 1+0 is used required disk capacity is doubled to accommodate disk mirroring (shadowing). Where RAID 5 is used it is assumed that roughly 20 percent of disk capacity is consumed by parity data, which is used to recover from failures. As a result, required disk space is increased by 25 percent.

[C] RAID configurations are expressed as AxDxC, where A equals number of storage array subsystems, D equals number disks within each subsystem, and C equals the capacity, in gigabytes, of each disk within the subsystem. Where "0" storage array subsystems are denoted the disk storage is provided by a single disk.

5.4.4 Distributed Process/Replicated Data for Publication Architecture

This architecture includes centralized transaction processing data management resources. However, unlike architectures that employ a fully centralized data management strategy, this architecture allows for the distribution of resources associated with decision support data management and with application and presentation services.

With primary-site replication for data publication, the primary or master data resource copies data to multiple target data stores. However, the replicated data is changed only at the primary site. The most simplistic example of this model is a single primary site that replicates all its data to a target system or set of identical systems. In another, more complicated, configuration, a primary-site DBMS could also be partitioned to multiple target sites via replication. In this situation, replication would copy specific portions of the primary database to certain target sites. This data replication configuration is often used to create hot-site backups and to populate distributed decision support system data repositories.

This architecture is typically characterized as:

- **Scaleable** – Common processing resources are not required to simultaneously support data management operations associated with decision support and transaction processing activities. As a result, specialized and heterogeneous resources can be employed to deliver required services.
- **Loosely coupled** – Autonomous operations do not rely on a single, centralized data resource for decision support and transaction processing activities. Rather, disparate activities, such as those associated with decision support and transaction processing, leverage specialized and largely independent resources.
- **Complex** – Disparate technologies mandate specialized skills and increase complexity of management, operation, and administration of heterogeneous technologies.
- **Fault tolerant** – Distributed data management and application processing allow for increased fault tolerance, as all system services and associated processing do not rely on any single resource.
- **Specialized** – Provides opportunity to use specialized technology resources, which are specifically engineered for supporting and providing particular services.
- **Flexible** –System growth, and service support not constrained by single hardware environment and limited vendor or third party product offerings.

With these characteristics in mind, the architecture for Project EASI/ED was been designed to mitigate weaknesses and leverage strengths intrinsic to centralized process/replication for data publication candidate framework architecture. This architecture is detailed in Figure 5.20, shown on the following page.

Figure 5.20 describes interfaces between hardware components and complements the software interface descriptions provided in Appendix C.

The distributed process/replicated data for publication candidate framework architecture is organized in terms of five technical "classifications." These classifications are:

- **Electronic Commerce Resources** – provide the software and hardware technologies required to support telephony, Web, and EDI-based processing. Specifically, these technologies allow users to access and invoke system services, such as those provided by transaction processing and decision support technologies, through commonly available Web browsers and telephone technology. Additionally, these technologies allow trading partners to electronically communicate via X12 compliant transaction sets.

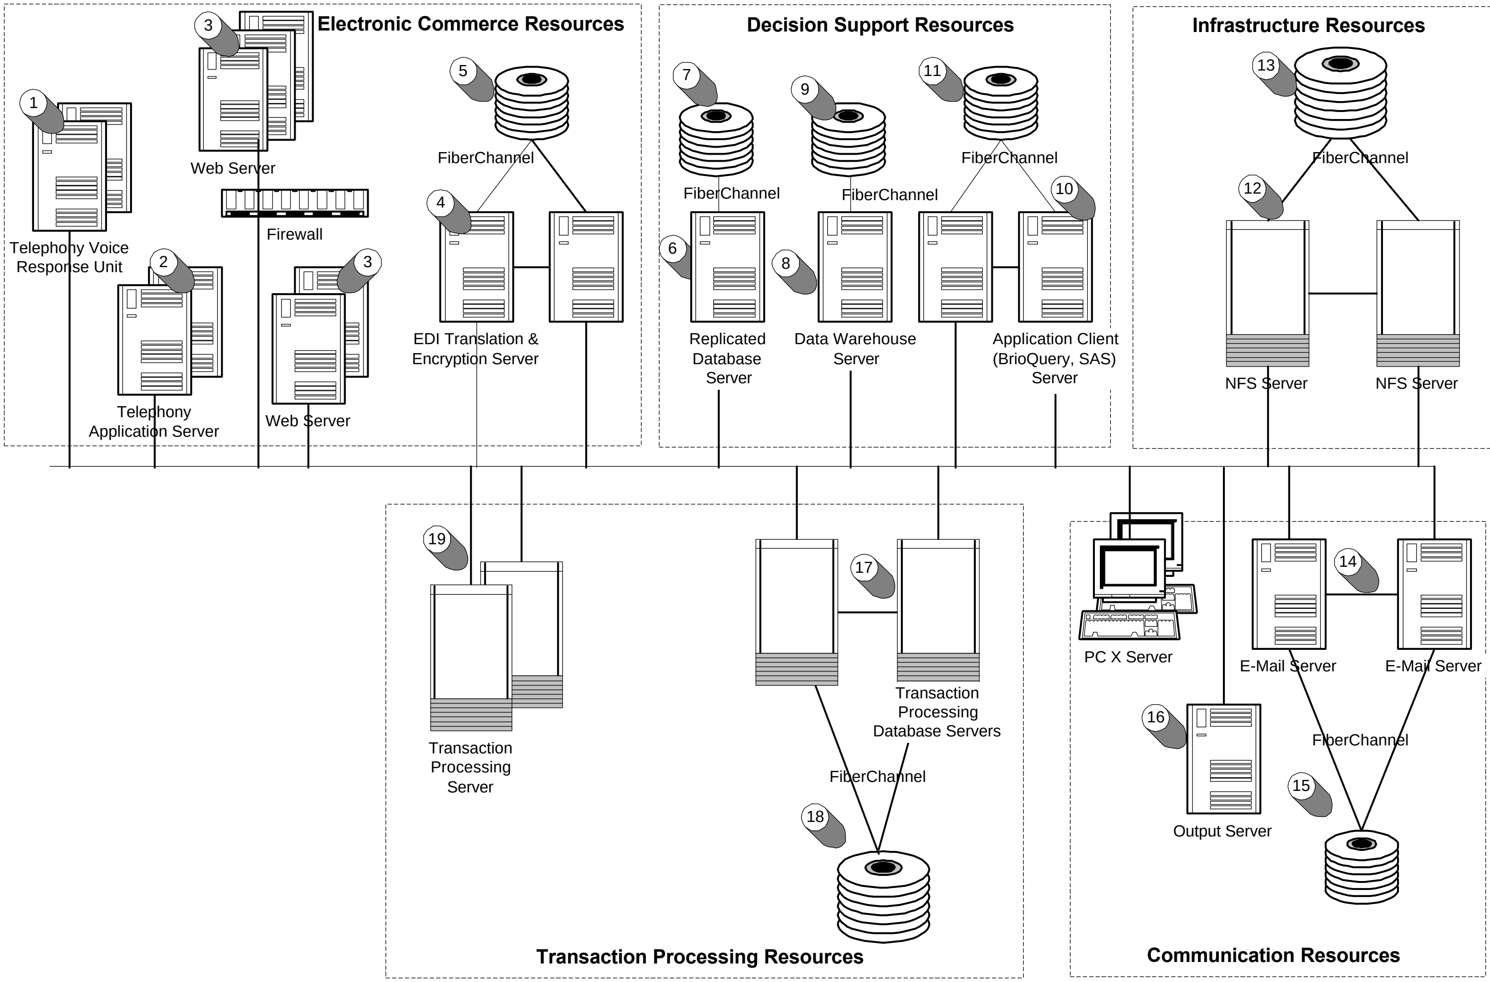

Figure 5-20. Distributed Process/ Replication for Data Publication Architecture

- **Infrastructure Resources** – provide the software and hardware technologies required to support the "Centralized Software Management and Dynamic Deployment" strategy described in Subsection 5.3.7. These technologies support DNS name services and provide the NFS services required to support the AutoClient server configurations that are used for the Web, transaction processing, and application display servers. These services are largely transparent to the user, but simplify system administration, maintenance, and operations.

- **Communication Resources** – provide the software and hardware technologies required to support E-mail, printer, and facsimile services. Additionally, these technologies provide the "client-side" remote application display services that is used to invoke and execute system services via LAN/WAN-based X Window System clients.

- **Transaction Processing Resources** – provide the software and hardware technologies required to support on-line and batch-oriented transaction processing operations. Transaction processing servers will be used to provide "component-based" transaction processing services. These services will be invoked in batch, via electronic commerce technologies (telephony and Web-based), and by communication resources (X Window System clients).

  Unlike previously described architectures, a single cluster of database servers is exclusively dedicated to supporting transaction-processing operations. Decision support operations are executed against other resources and are not supported by the transaction processing database servers.

- **Decision Support Resources** – provide the software and hardware technologies required by decision support operations. One server is dedicated to data warehousing operations, while a second server supports a "read-only" replica subset of the transaction processing database. The warehouse is loaded from the transaction processing databases and includes derived and aggregated data.

  Application display servers are required to support the "Centralized Execution and Remote Display of Application Clients" – a strategy described in Subsection 5.3.8. Application display servers support execution of multiple client application software instantiations, which are remotely displayed on remote X Window System communication resources. These clients primarily include decision support tools such as BrioQuery and SAS. However, X Window System transaction process application clients also leverage this server's resources.

Figure 5-21, describes the hardware and software components that comprise the distributed process/replicated data for publication candidate framework architecture.

Reference	Vendor/Product	Description	Quantity	Hardware Specification
1	VoiceTek, Inc VTK 3000	Interactive Voice Response Unit Server	3	- 1 Pentium processor - 166Mhz clock speed - 32MB RAM - 4 GB of Disk - SCO UNIX 5.0.2 - T1 trunk interface - 100BaseT LAN interface
	VoiceTek, Inc. Generations Telephony Server Platform v4.0	Telephony – Interactive Voice Response Unit Software	3	
2	Sun Microsystems, Inc. Ultra Enterprise 1	Enterprise Class Telephony Application Hardware Server	6	- 1 UltraSPARC processor - 167Mhz clock speed - 64 MB RAM - 2 GB Disk - Solaris 2.6 - 100BaseT LAN Interface
	VoiceTek, Inc. Generations Runtime Server Platform v4.0	Telephony – Runtime Application Server	6	
3	Sun Microsystems, Inc. Ultra Enterprise 1	Web Hardware Server	3	- 1 UltraSPARC processor - 167Mhz clock speed - 64 MB RAM - 2 GB Disk - Solaris 2.6 - 100BaseT LAN Interface
4	Sun Microsystems, Inc Ultra Enterprise 2	Highly Available, Clustered Enterprise Class EDI Hardware Server	2	- 1 UltraSPARC processor - 167 Mhz clock speed - 256 MB RAM - Solaris 2.6 - 100BaseT LAN Interface
	Solstice HA v1.3	High Availability Hardware Monitoring and Fault Detection Software	1	
	Solstice: DiskSuite v4.0	High Availability Volume Management Software	1	
	Premenos Corporation Templar Enterprise	EDI Security and SMTP/MIME Delivery Software	1	
	Premenos Corporation EDI/Open	EDI X12 Translation Software	1	
	Oracle Corporation Oracle 8 Server	Relational Database Management System Software	1	
5	SPARCstorage Array 112 RSM	Storage Array Subsystem	1	- 1x110 Mhz MicroSPARC II controller - 6 SCSI-2 buses - 13x2 GB Disks (26 GB total) - 2x25 Mbps FibreChannel (FC) Interfaces - Configured for RAID 5
6	Sun Microsystems, Inc. Ultra Enterprise 3000	Enterprise Class Decision Support Database Hardware Server	1	- 2 UltraSPARC processors - 250 Mhz clock speed - 1 GB RAM - Solaris 2.6 - 100BaseT LAN Interface

Figure 5-21. Distributed Process/Replicated Data Publication Architecture Components

Reference	Vendor/Product	Description	Quantity	Hardware Specification
6	Solstice: DiskSuite v4.0	High Availability Volume Management Software	1	-
	Oracle Corporation Oracle 8 Server	Relational Database Management System Software	1	-
7	SPARCstorage Array 214 RSM	Storage Array Subsystem	9	- 1x110 Mhz MicroSPARC II controller - 6 SCSI-2 buses - 42x4 GB Disks (2184 GB total) - 2x25 Mbps FibreChannel (FC) Interfaces - Configured for RAID 5
8	Sun Microsystems, Inc. Ultra Enterprise 3000	Enterprise Class Decision Support Database Hardware Server	1	- 3 UltraSPARC processors - 250 Mhz clock speed - 2 GB RAM - Solaris 2.6 - 100BaseT LAN Interface
	Solstice: DiskSuite v4.0	High Availability Volume Management Software	1	
	Red Brick Systems, Inc. Red Brick Warehouse v5.0	Data Warehouse Relation al Database Management system Software	1	
9	SPARCstorage Array 214 RSM	Storage Array Subsystem	9	- 1x110 Mhz MicroSPARC II controller - 6 SCSI-2 buses - 42x4 GB Disks (2520GB total) - 2x25 Mbps FibreChannel (FC) Interfaces - Configured for RAID 5
10	Sun Microsystems, Inc. Ultra Enterprise 2	High Availability Cluster Enterprise Class Application Display Hardware Server	2	- 2 UltraSPARC processors - 300 Mhz clock speed - 1 GB RAM - Solaris 2.6 - 100BaseT LAN Interface
11	SPARCstorage Array 112 RSM	Storage Array Subsystem	2	- 1x110 Mhz MicroSPARC II controller - 6 SCSI-2 buses - 19x2 GB Disks (76GB total) - 2x25 Mbps FibreChannel (FC) Interfaces - Configured for RAID 5
12	Sun Microsystems, Inc. Ultra Enterprise 2	Highly Available, Clustered Enterprise Class NFS Hardware Servers	2	- 1 UltraSPARC processor - 167 Mhz clock speed - 128 MB RAM - Solaris 2.6 - 100BaseT LAN Interface
	Netscape Communication Corporation Enterprise Server v3.0	Web Server Software	3	
	Brio Technology, Inc. BrioQuery.Insight	Online Analytical Processing Plug-in for the Netscape Enterprise Server	3	

Figure 5-21. Distributed Process/Replicated Data Publication Architecture Components (cont'd)

Reference	Vendor/Product	Description	Quantity	Hardware Specification
12	Brio Technology, Inc. BrioQuery.Quickview	Online Analytical Processing Plug-in for the Netscape Enterprise Server	3	
	Netscape Communication Corporation Directory Server	Web Directory Server Software	1	
	Brio Technology, Inc. BrioServer	Data Warehouse Extract Utility for Online Analytical Processing	1 (site license)	
	SAS Institute, Inc SAS	Statistical and Quantitative Analysis Software	1 (site license)	
	Solstice HA v1.3	High Availability Hardware Monitoring and Fault Detection Software	1	
	Solstice: DiskSuite v4.0	High Availability Volume Management Software	1	
	HA NFS	Highly Available Network File System Software	1	
	RRDNS	Primary Round-Robin Domain Name Server Software	1	
13	SPARCstorage Array 112 RSM	Storage Array Subsystem	2	- 1x110 Mhz MicroSPARC II controller - 6 SCSI-2 buses - 19x2 GB Disks (76 GB total) - 2x25 Mbps FibreChannel (FC) Interfaces - Configured for RAID 5
14	Sun Microsystems, Inc. Ultra Enterprise 2	Highly Available, Clustered Enterprise Class Electronic Mail Hardware Servers	2	- 1 UltraSPARC processor - 167 Mhz clock speed - 450 MB RAM - Solaris 2.6 - 100BaseT LAN Interface
	Sun Microsystems, Inc. Solstice Internet Mail Server v2.0	Internet Electronic Mail Server Software	1	
	Solstice HA v1.3	High Availability Hardware Monitoring and Fault Detection Software	1	
	Solstice: DiskSuite v4.0	High Availability Volume Management Software	1	
	RRDNS	Secondary Round-Robin Domain Name Server Software	1	
15	SPARCstorage Array 219 RSM	Storage Array Subsystem	3	- 1x110 Mhz MicroSPARC II controller - 6 SCSI-2 buses - 18x9 GB Disks (486 GB total) - 2x25 Mbps FibreChannel (FC) Interfaces - Configured for RAID 10

Figure 5-21. Distributed Process/Replicated Data Publication Architecture Components (cont'd)

Reference	Vendor/Product	Description	Quantity	Hardware Specification
16	Sun Microsystems, Inc. Ultra Enterprise 1	Output Management Hardware Server	1	- 1 UltraSPARC processor - 167Mhz clock speed - 64 MB RAM - 4 GB Disk - Solaris 2.6 - 100BaseT LAN Interface
	DAZEL Corporation DAZEL Output Server	Output Management System Software	1	
17	Sun Microsystems, Inc. Ultra Enterprise 6000	Highly Available, Clustered Enterprise Class Database Management Hardware Servers	2	- 18 UltraSPARC processors - 250 Mhz clock speed - 3 GB RAM - Solaris 2.6 - 100BaseT LAN Interface
	Oracle Corporation Oracle 8 Parallel Server	Parallel Processing Relational Database Management System Software	1	
	Solstice PDB v1.2	High Availability, Parallel Processing Hardware Monitoring and Fault Detection Software	1	
	Solstice: DiskSuite v4.0	High Availability Volume Management Software	1	
18	SPARCstorage Array 214 RSM	Storage Array Subsystem	36	- 1x110 Mhz MicroSPARC II controller - 6 SCSI-2 buses - 42x4 GB Disks (7392 GB total) - 2x25 Mbps FibreChannel (FC) Interfaces - Configured for RAID 10
19	Sun Microsystems, Inc Ultra Enterprise 1	Batch and Online Transaction Processing Monitor Software	15	- 1 UltraSPARC processor - 167 Mhz clock speed - 512 MB RAM - 2.1 GB Disk - Solaris 2.6 - 100BaseT LAN Interface
	BEA Systems, Inc. BEA Tuxedo v6.x	Online Transaction Processing Monitor	15	
	BEA Systems, Inc. HA Extensions – HA Client Wrapper	High Availability, In-flight transaction recovery and resubmission software	15	
	BEA Systems, Inc. BEA Jolt v1.1	Web Gateway for Online Transaction Processing Monitor Services	15	

Figure 5-21. Distributed Process/Replicated Data Publication Architecture Components (cont'd)

The following table, Figure 5-22, details storage capacity calculations for the distributed process/replication for data publication candidate architecture. These calculations were used to estimate required storage capacity.

Service	Raw Data Calculation[B]	Total Capacity	RAID Configuration	RAID Calculation[C]	Instances	Total RAID Capacity	Disk Configuration[A]
Telephony Service	2GB	2 GB	NA	NA	1	12GB	0x1x2
Web Service	2GB	2GB	NA	NA	3	6GB	0x1x2
EDI Service	20GB	20GB	RAID 5	25% increase	1	25GB	1x13x2
Decision Support RDBMS Service	590GB	1180GB	RAID 5	25% increase	1	1475GB	9x42x4
Decision Support Data Warehouse Service	730GB	1460GB	RAID 5	25% increase	1	1505GB	9x42x4
NFS Service	60GB	60GB	RAID 5	25% increase	1	76GB	2x19x2
Application Service	60GB	60GB	RAID 5	25% increase	1	76GB	2x19x2
Electronic Mail Service	378GB	378GB	RAID 5	25% increase	1	478GB	3x18x9
Output Management Service	4GB	4GB	NA	NA	1	4GB	0x1x2
OLTP Database Service	1490GB	2980GB	RAID 10	100% increase	1	5960GB	36x42x4
OLTP Monitor Service	2GB	2GB	NA	NA	15	30GB	0x1x2
Total Required Storage Capacity Requirements						9647GB	

Figure 5-22. Distributed Process/ Replication for Data Publication Architecture Storage Capacity

[A] RAID configurations are expressed as AxDxC, where A equals number of storage array subsystems, D equals number disks within each subsystem, and C equals the capacity, in gigabytes, of each disks within the subsystem. Where "0" storage array subsystems are denoted the disk storage is provided by a single disk.

[B] According to the reference text *Capacity and Configuration Planning for Solaris Servers*, unless more accurate information is available, DBMS' should be configured with approximately twice as much disk space as raw data to account for additional database structures, such as table indexes, temporary tables, log files, archives, etc. Taking this into account, the "Raw Data Calculation" represents the disk space available for raw data, assuming the remaining available disk space will be consumed by other database structures.

[C] Where RAID 1+0 is used, required disk capacity is doubled to accommodate disk mirroring (shadowing). Where RAID 5 is used it is assumed that roughly 20 percent of disk capacity is consumed by parity data, which is used to recover from failures. As a result, required disk space is increased by 25 percent.